



US010528607B2

(12) **United States Patent**  
**Dinga et al.**

(10) **Patent No.:** **US 10,528,607 B2**  
(45) **Date of Patent:** **Jan. 7, 2020**

- (54) **SYNTAX TEMPLATES FOR CODING**
- (71) Applicant: **SPLUNK INC.**, San Francisco, CA (US)
- (72) Inventors: **Jindrich Dinga**, San Francisco, CA (US); **Yuan Xie**, San Francisco, CA (US); **Katherine Kyle Feeney**, San Francisco, CA (US); **Jesse Miller**, Berkeley, CA (US)
- (73) Assignee: **SPLUNK INC.**, San Francisco, CA (US)

6,427,142 B1 \* 7/2002 Zachary ..... G06F 8/34  
706/11  
6,658,646 B1 \* 12/2003 Hernandez, III ... G06F 9/45512  
717/114  
6,931,622 B1 \* 8/2005 Aldrich ..... G06N 5/025  
706/45

(Continued)

**OTHER PUBLICATIONS**

Title: The Cornell program synthesizer: a syntax-directed programming environment author: T Teitelbaum, published on 1981.\*

(Continued)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 46 days.

*Primary Examiner* — Chameli Das

(74) *Attorney, Agent, or Firm* — Shook, Hardy & Bacon, L.L.P.

(21) Appl. No.: **15/223,598**

(22) Filed: **Jul. 29, 2016**

(65) **Prior Publication Data**

US 2018/0032316 A1 Feb. 1, 2018

(51) **Int. Cl.**

**G06F 9/44** (2018.01)  
**G06F 16/33** (2019.01)  
**G06F 8/33** (2018.01)  
**G06F 11/30** (2006.01)

(52) **U.S. Cl.**

CPC ..... **G06F 16/334** (2019.01); **G06F 8/33** (2013.01); **G06F 11/30** (2013.01)

(58) **Field of Classification Search**

CPC ..... G06F 8/33; G06F 11/30; G06F 16/334  
See application file for complete search history.

(56) **References Cited**

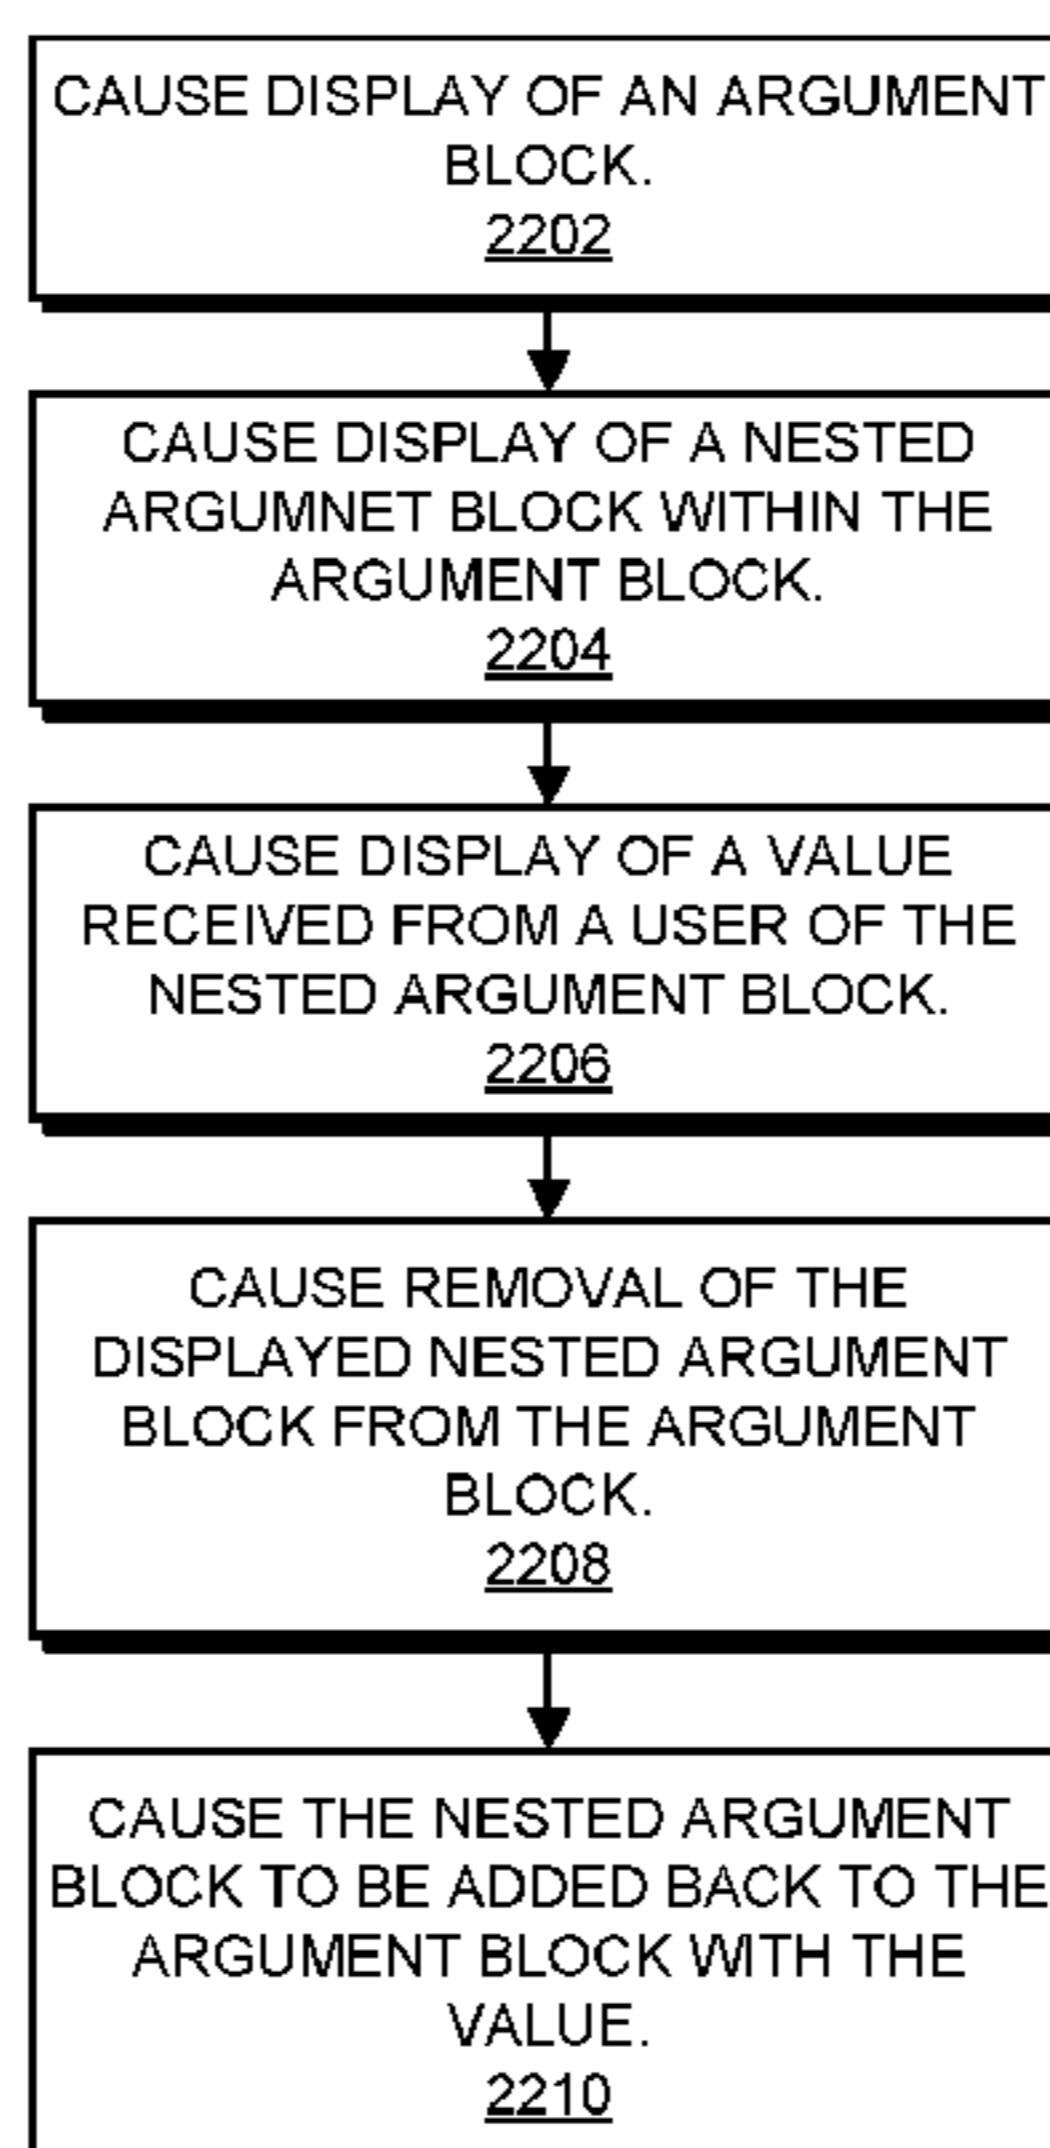
**U.S. PATENT DOCUMENTS**

5,603,021 A \* 2/1997 Spencer ..... G06F 17/246  
5,960,437 A \* 9/1999 Krawchuk ..... G06F 16/9024  
707/702

(57) **ABSTRACT**

Various approaches for automating code completion are described herein. More particularly, approaches are provided that automatically generate coded commands of a coding language (i.e., code) that function and operate as intended by the user. As the user codes the commands, such approaches assist a user in various ways. For example, such automated assistance provides the user an understanding of various coding options available in the coding language. The assistance also enforces the proper employment of the available coding options, as well as provides an understanding of the functionality of the generated code. Automating code completion provides various benefits to the user, such as decreasing the time the user spends generating code, increasing the likelihood that the generated code functions and operates on a system as intended, and reducing the number of code versions required to be executed or compiled by the system.

**30 Claims, 30 Drawing Sheets**



(56)

References Cited

U.S. PATENT DOCUMENTS

7,360,202 B1 \* 4/2008 Seshadri ..... H04L 12/1859  
717/106  
8,631,323 B1 \* 1/2014 Gray ..... G06F 8/33  
715/255  
9,483,459 B1 \* 11/2016 Riley ..... G10L 25/48  
2003/0028864 A1 \* 2/2003 Bowen ..... G06F 8/48  
717/141  
2003/0105620 A1 \* 6/2003 Bowen ..... G06F 17/5022  
703/22  
2003/0120698 A1 \* 6/2003 Wyke ..... G06F 9/44526  
718/100  
2006/0075328 A1 \* 4/2006 Becker ..... G06F 17/246  
715/213  
2006/0282818 A1 \* 12/2006 DeSpain ..... G06F 17/246  
717/109  
2007/0288429 A1 \* 12/2007 Liu ..... G06F 17/30935  
2008/0172647 A1 \* 7/2008 Roche ..... G06F 17/30392  
717/100  
2010/0004980 A1 \* 1/2010 Bowen ..... G06Q 30/02  
705/7.31  
2012/0150939 A1 \* 6/2012 Zaifman ..... G06F 9/45512  
709/203  
2012/0191716 A1 \* 7/2012 Omoigui ..... H01L 27/1463  
707/740  
2012/0197928 A1 \* 8/2012 Zhang ..... G06F 17/30545  
707/769  
2013/0125094 A1 \* 5/2013 Wolfram ..... G06F 17/28  
717/109

2013/0159832 A1 \* 6/2013 Ingargiola ..... G06F 17/246  
715/220  
2013/0235197 A1 \* 9/2013 Ebrom ..... G06F 9/54  
348/143  
2014/0047413 A1 \* 2/2014 Sheive ..... H04L 65/403  
717/110  
2014/0189548 A1 \* 7/2014 Werner ..... G06F 3/0481  
715/762  
2015/0134633 A1 \* 5/2015 Colgrove ..... G06F 17/30424  
707/706  
2016/0188181 A1 \* 6/2016 Smith ..... G06F 3/048  
715/765  
2016/0225271 A1 \* 8/2016 Robichaud ..... G06F 17/30663  
2016/0357586 A1 \* 12/2016 Rose ..... G06F 9/445  
2018/0052879 A1 \* 2/2018 Wright ..... G06F 3/0605

OTHER PUBLICATIONS

Title: Syntax-directed editing, author: TF Lunney et al published on 1988.\*  
Work with HTML5 and CSS more efficiently. (Jun. 2014). Adobe Dreamweaver CC Tutorials. Retrieved from the Internet Nov. 4, 2016 at <<https://helpx.adobe.com/dreamweaver/how-to/using-code-view-dreamweaver.html>>. 41 pages.  
Using NetBeans IDE 5.5. (Oct. 2006). NetBeans. Retrieved from the Internet at <<http://netbeans.org>>. pp. 1-96.  
Use code completion to enter symbols. (2016). Xcode Help. Retrieved from the Internet Nov. 4, 2016 at <<http://help.apple.com/xcode/mac/8.0/#/dev8716704af>>. 2 pages.

\* cited by examiner

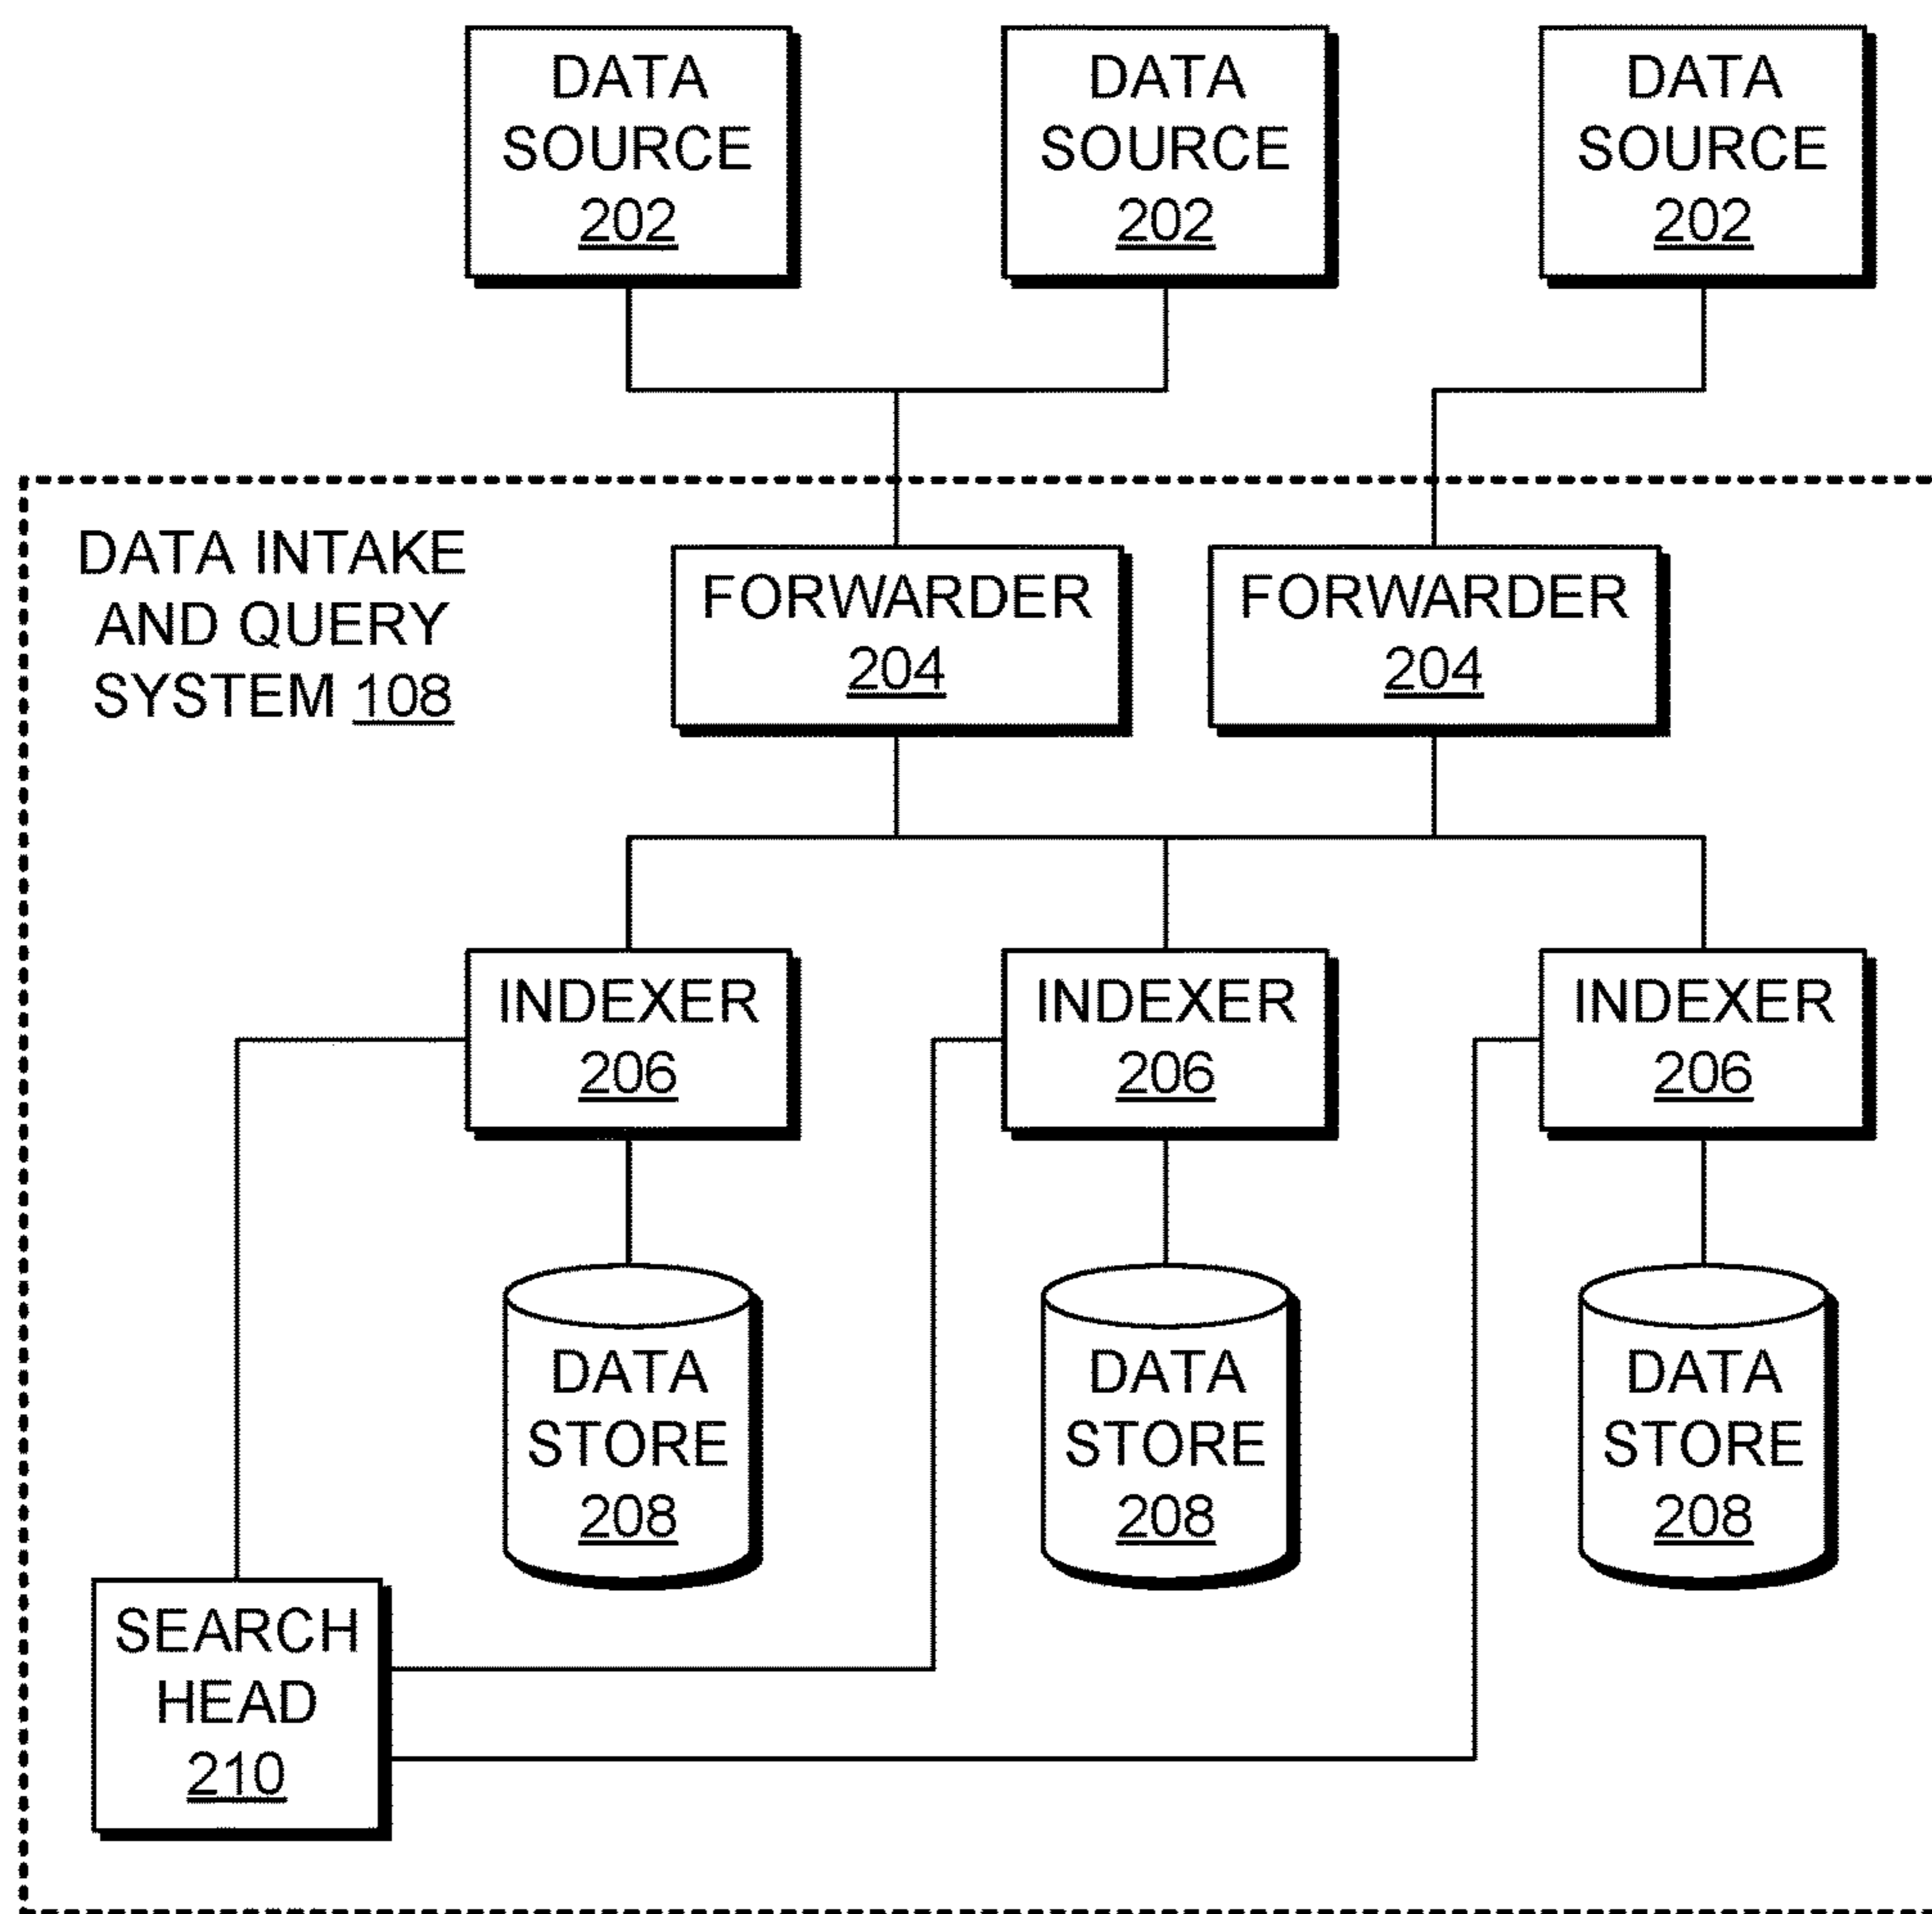
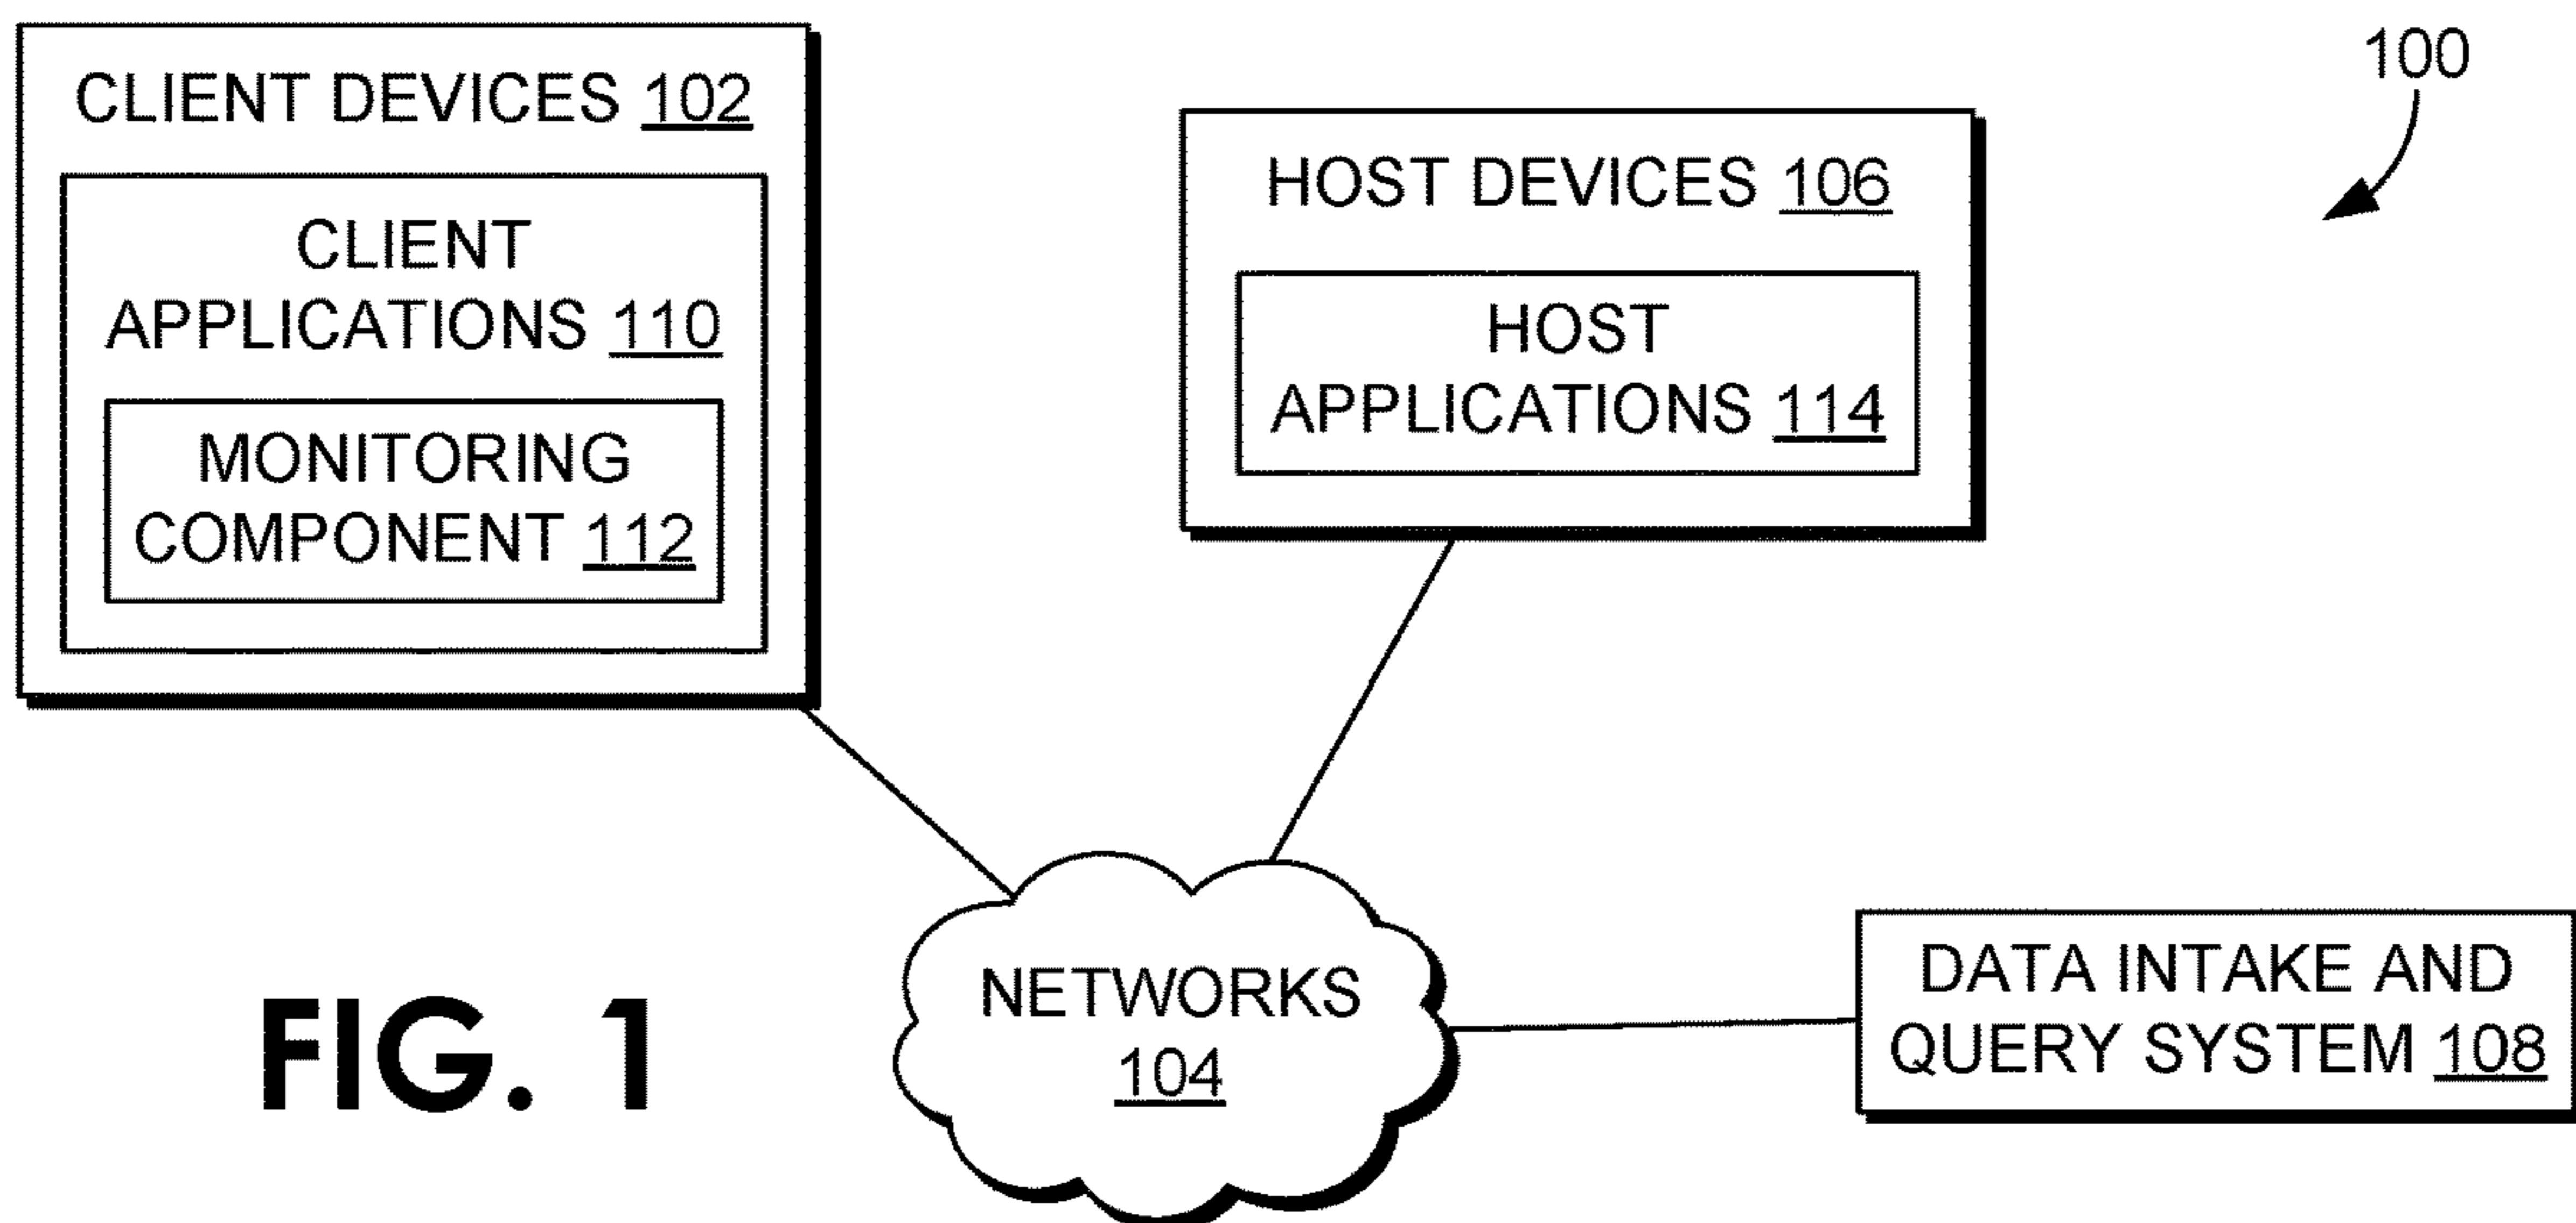
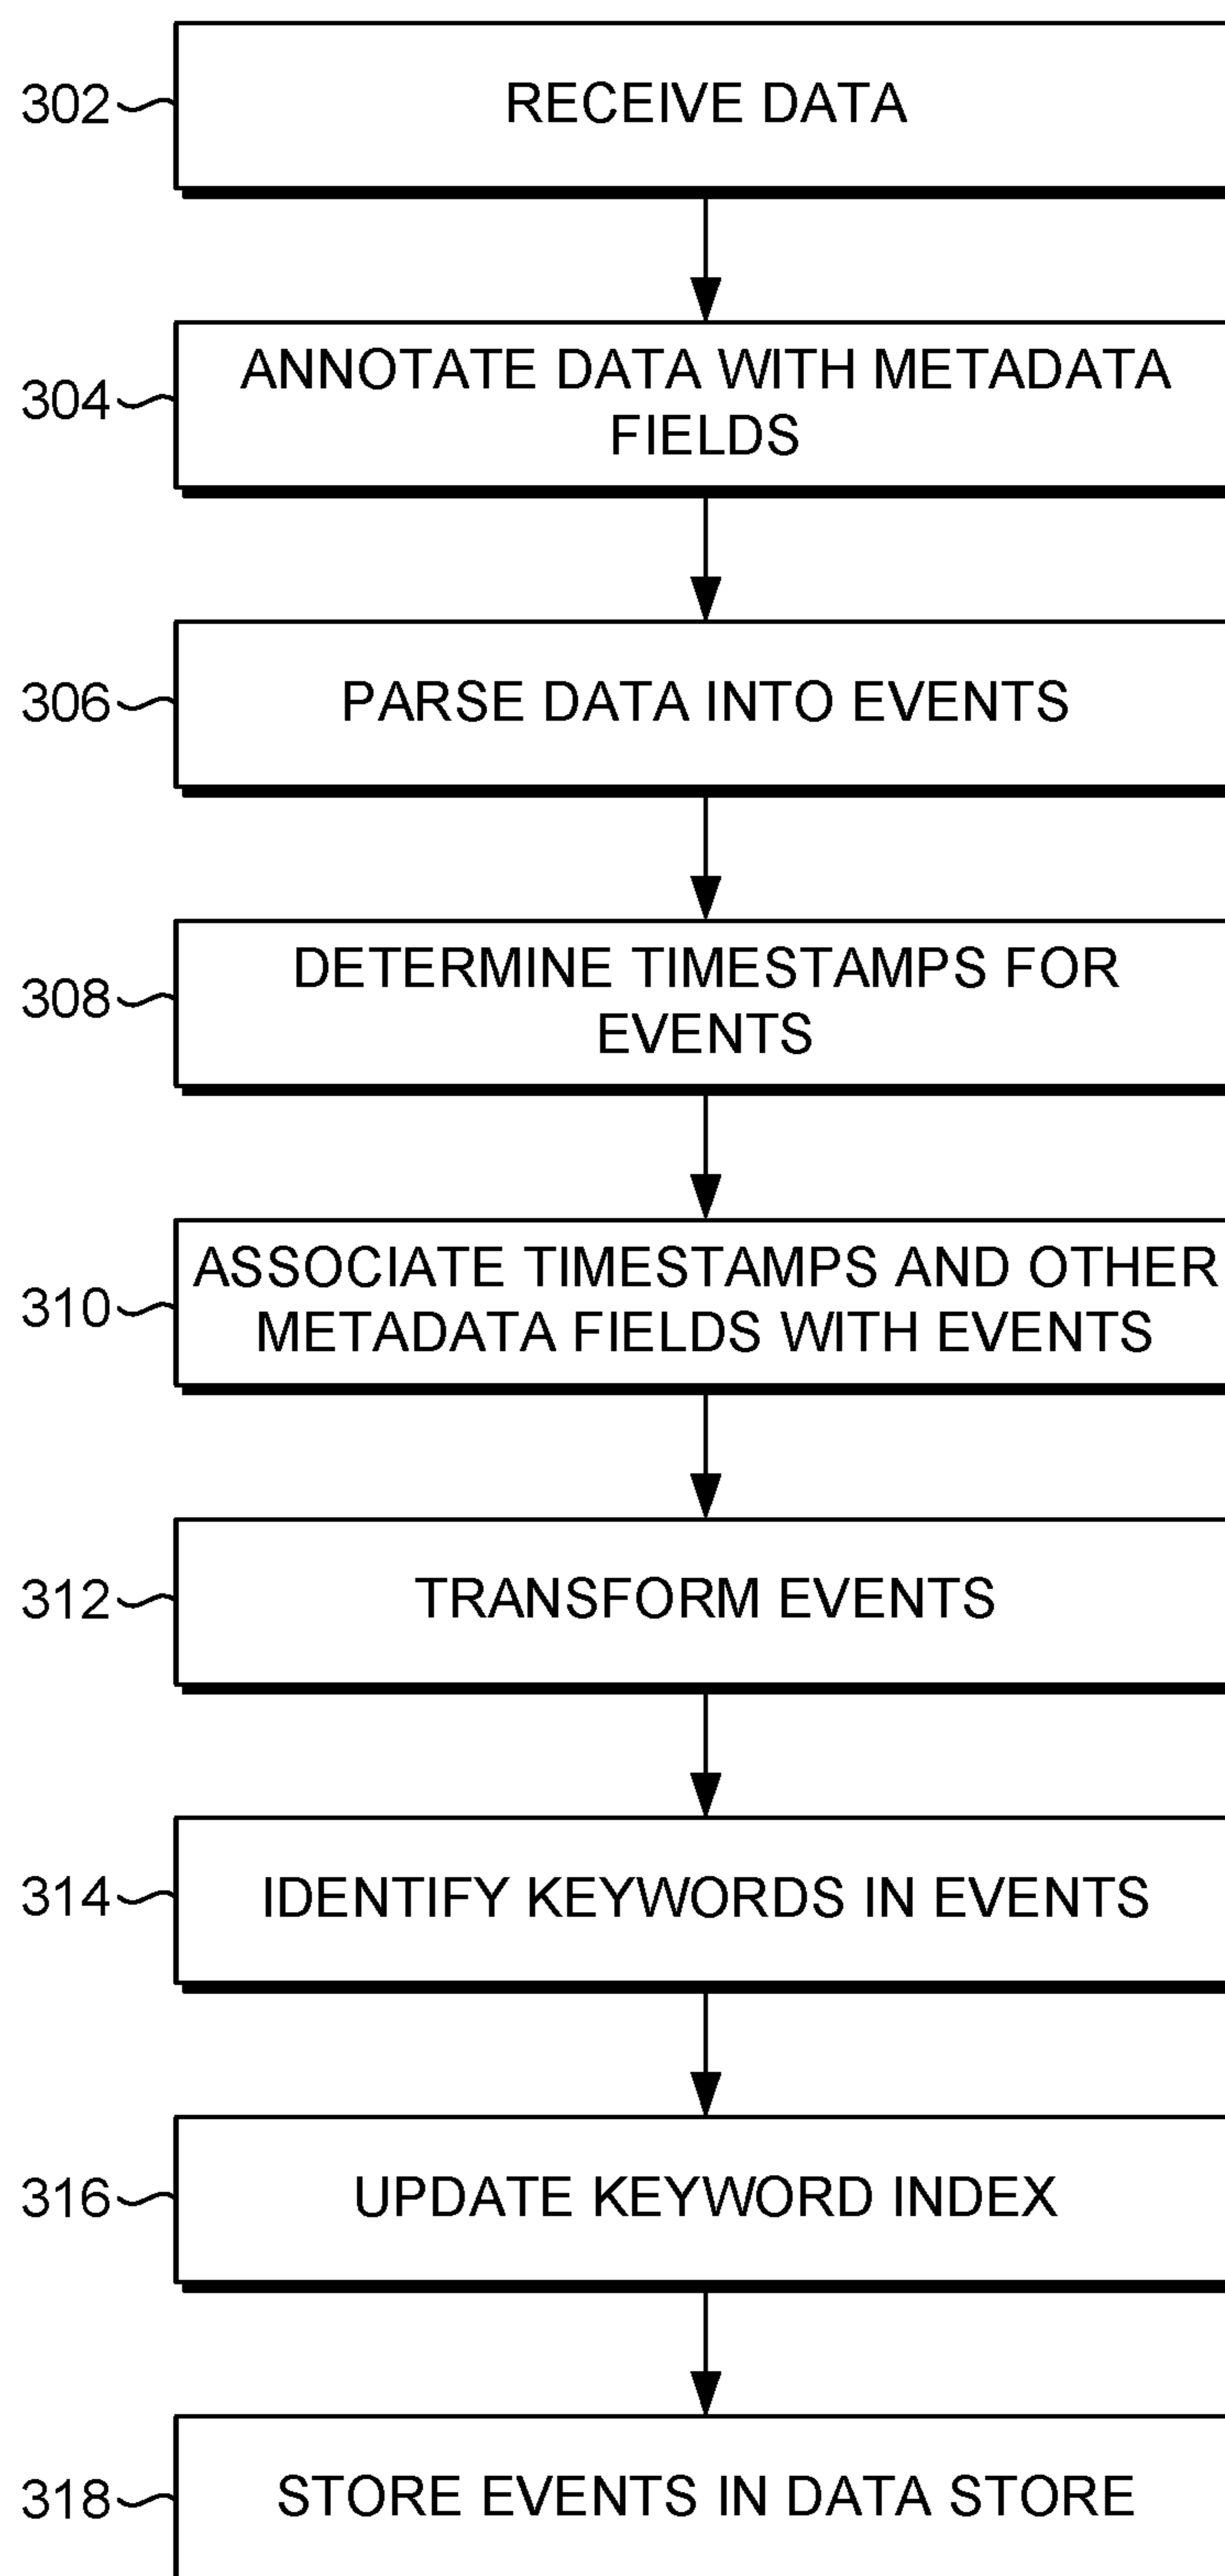
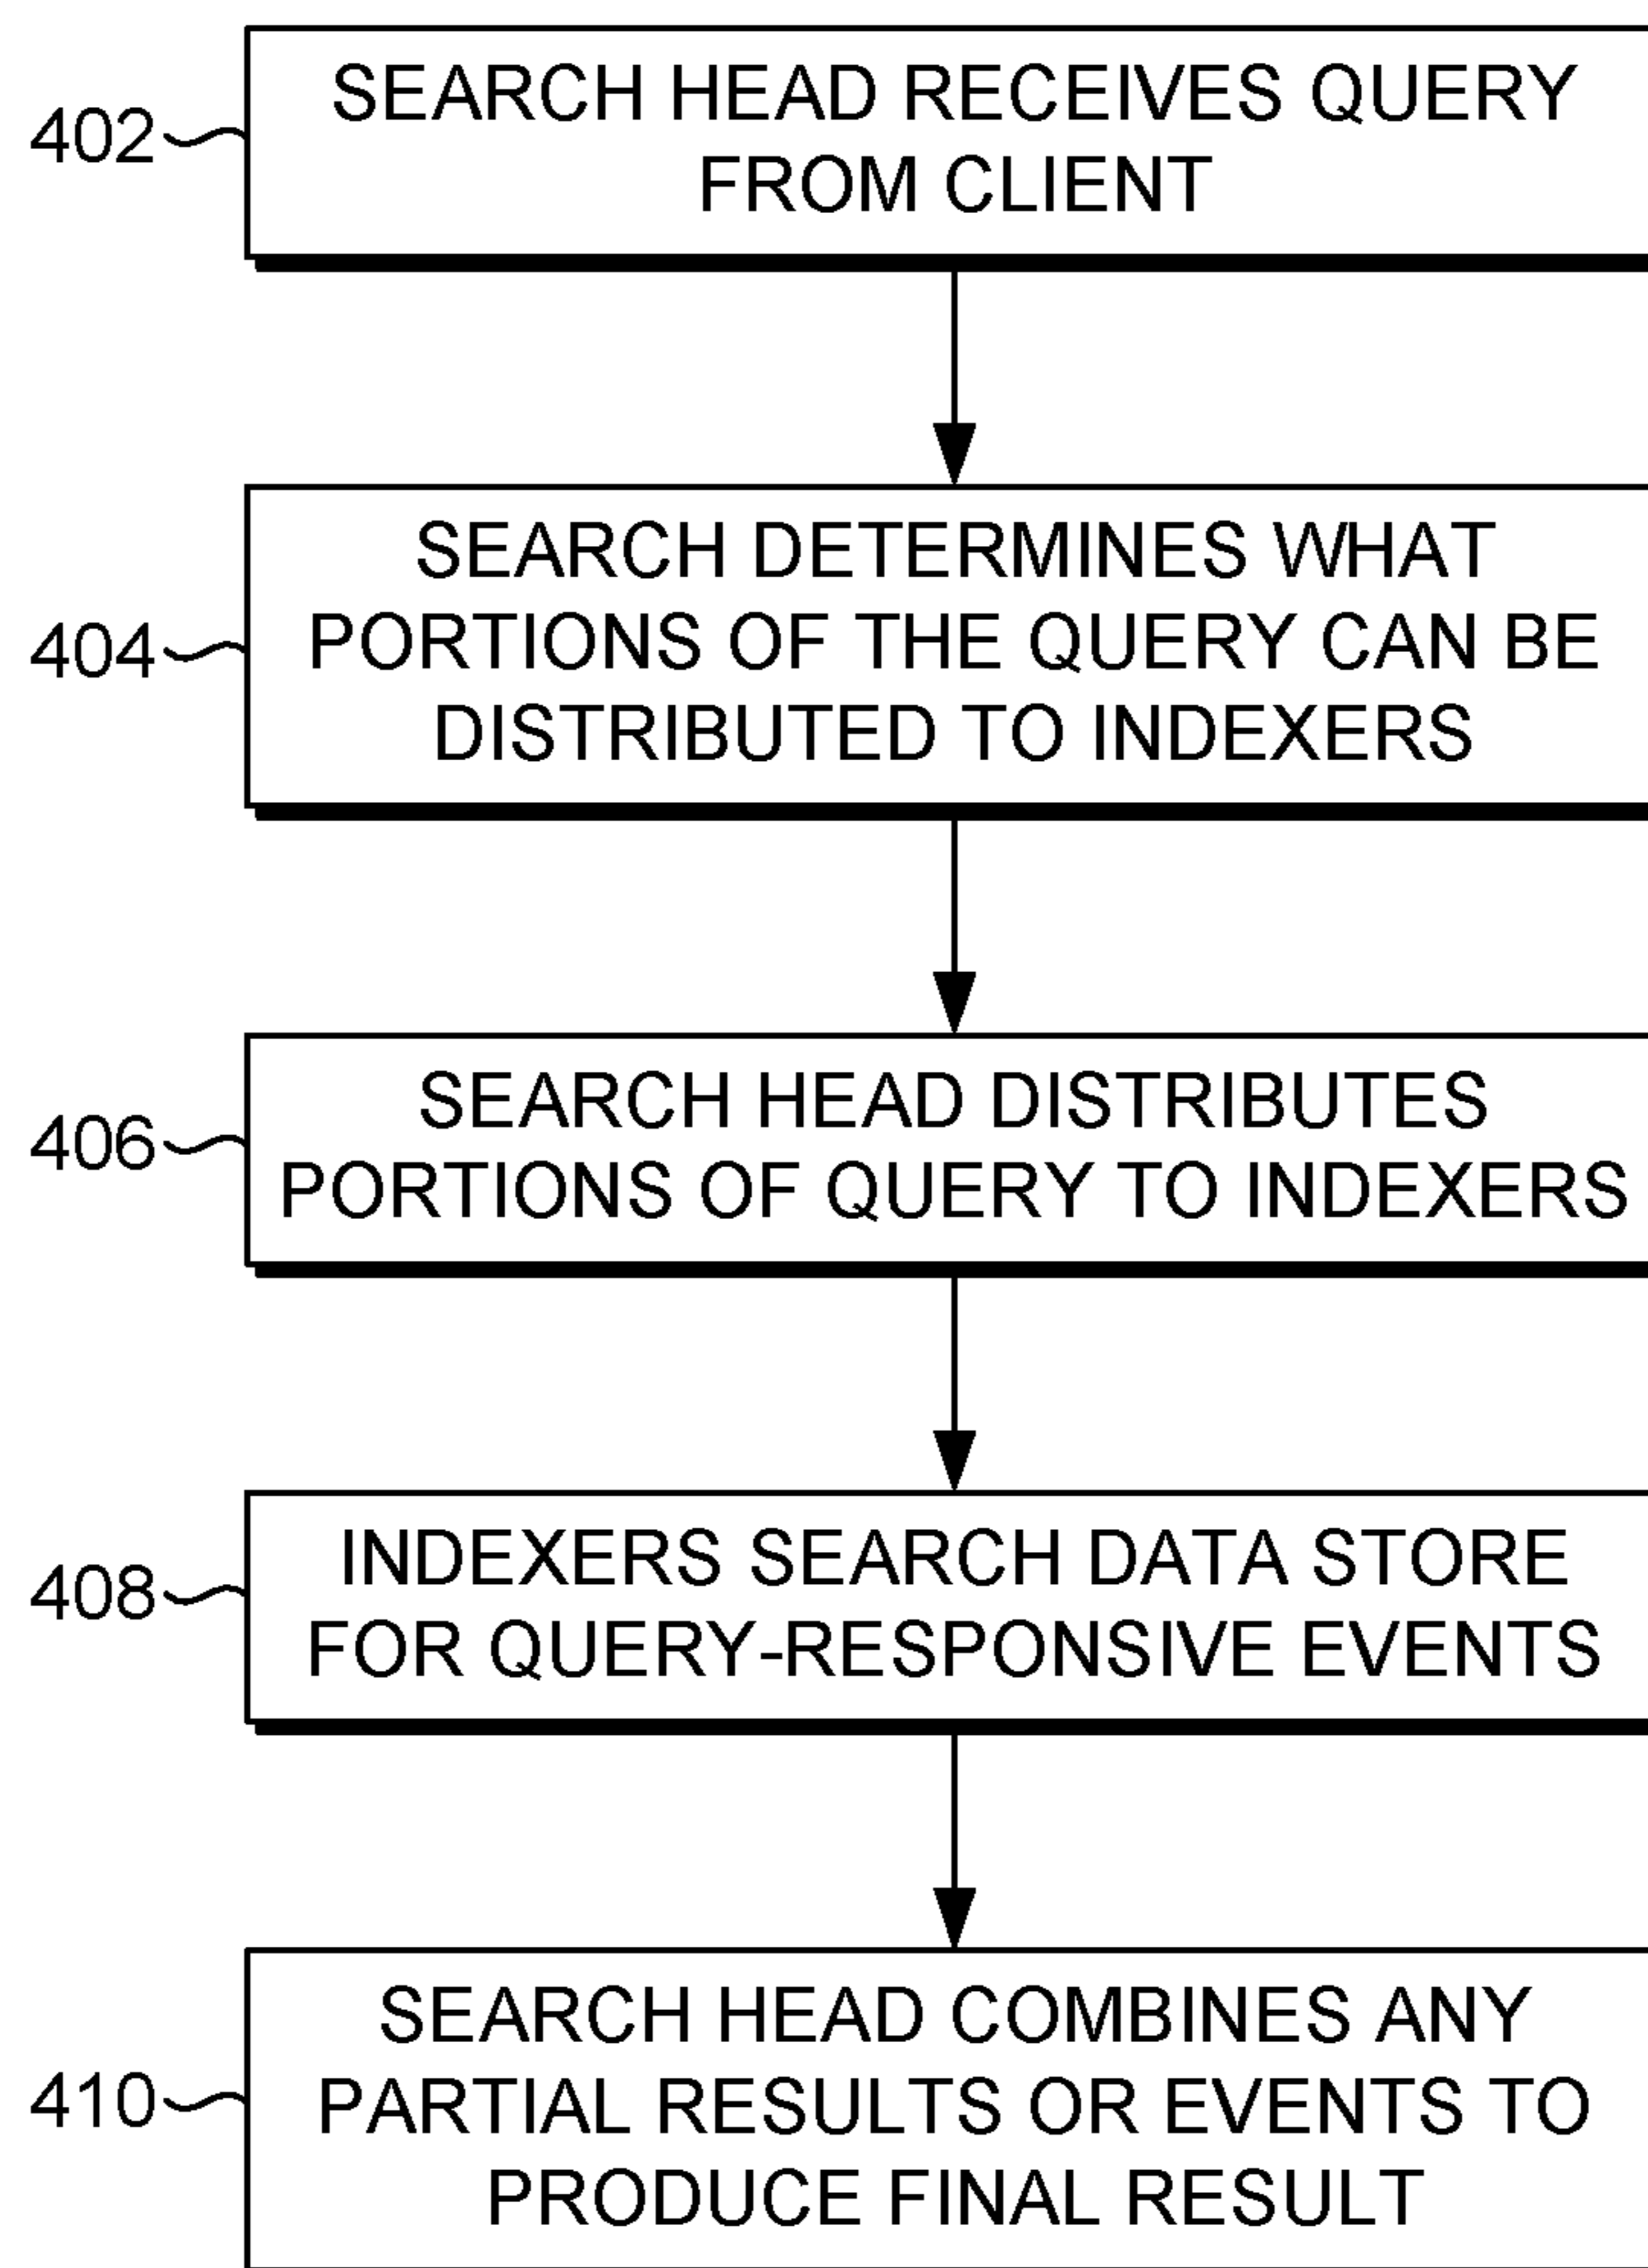


FIG. 2

**FIG. 3**

**FIG. 4**

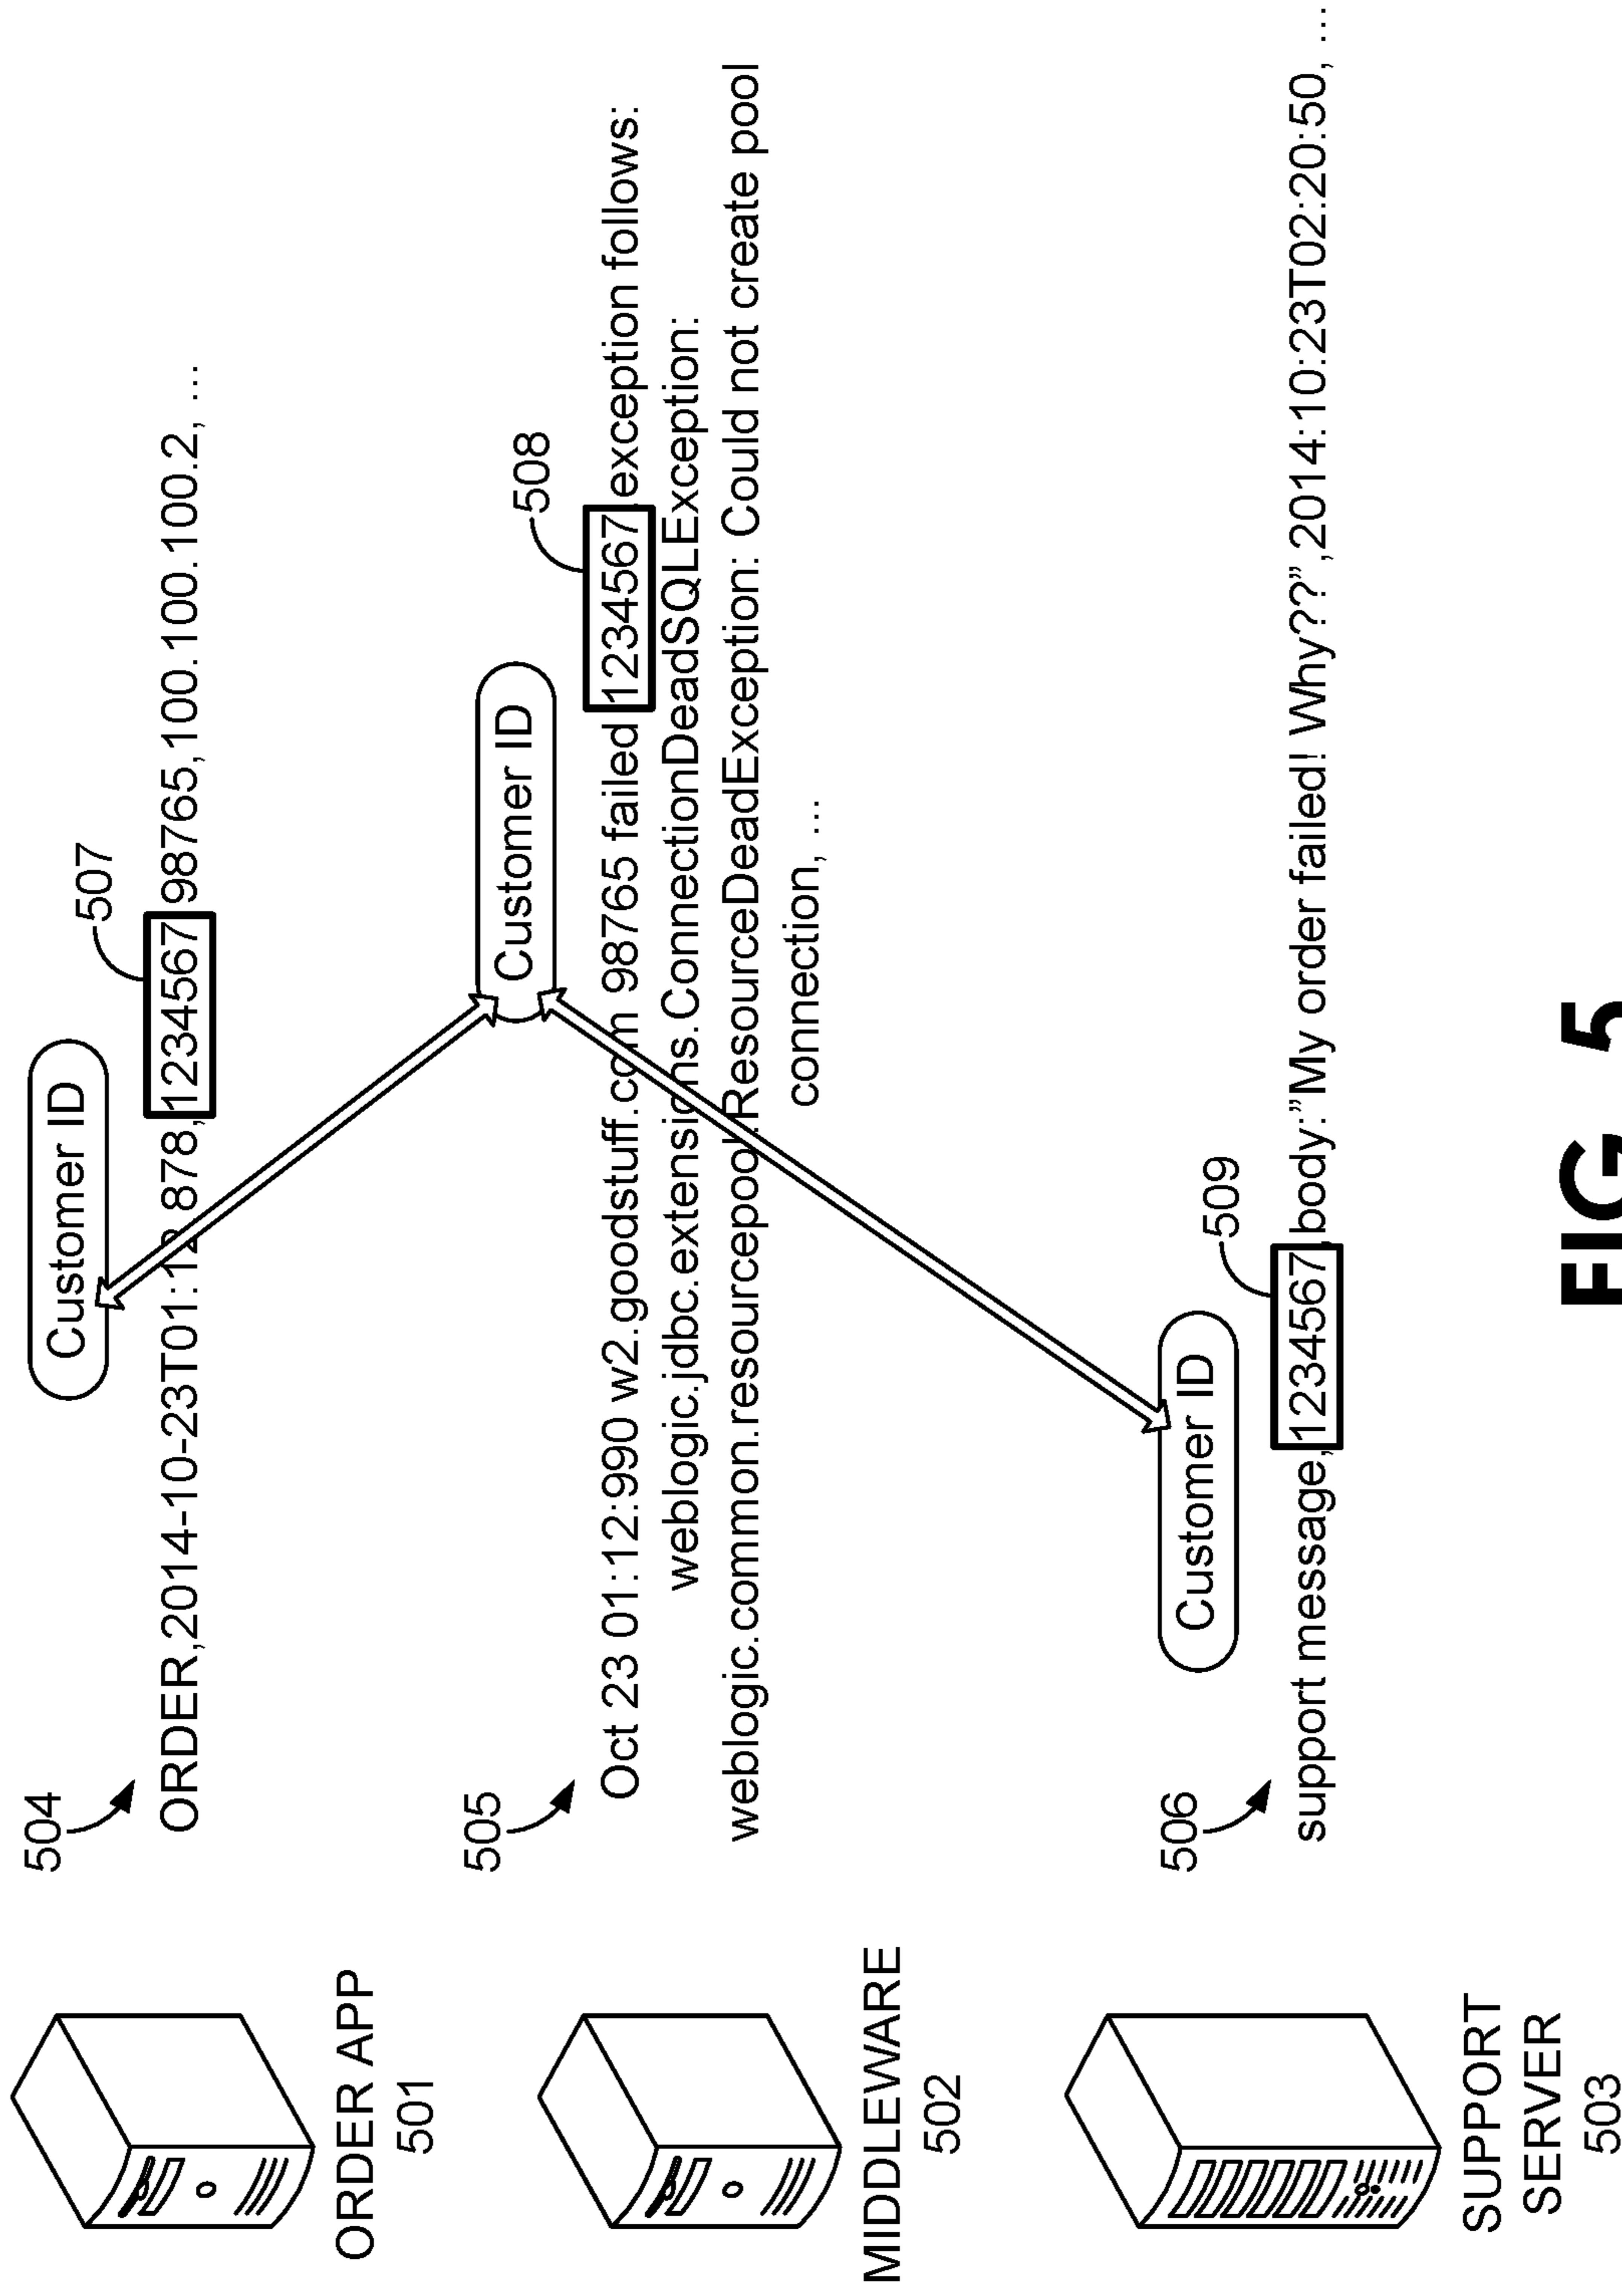


FIG. 5

Search Screen 600

Search Pivot Reports Alerts Dashboards Search & Reporting

Q New Search

buttercupgames Search Bar 602

✓ 36,819 events (before 4/30/14 2:19:02.000 PM)

Events (36,819) Statistics Visualization Search Results Tabs 604

Save as menu → Save As ▼ Close

Time Range Picker 612 All time ▼ Q

Job ▼ [Pause] [Stop] [Refresh] [Download] [Share] [Smart Mode ▼]

Search action buttons Search mode selector

Format Timeline ▼ - Zoom Out + Zoom to Selection x Deselect Timeline 605 1 hour per column

Timeline visualization

List ▼ Format ▼ 20 Per Page ▼ < Prev 1 2 3 4 5 6 7 8 9 ... Next >

Events List 608

i	Time	Event
>	4/28/14 6:22:16.000 PM	91.205.189.15 - - [28/Apr/2014:18:22:16] *GET /oldlink?itemId=EST-14&JSESSIONID=SD6SL7FF7ADFF53113 HTTP 1.1" 200 1665 "http://www.buttercupgames.com/oldlink?itemId=EST-14" Mozilla/5.0 (Windows NT 6.1; WOW 64) AppleWebKit/536.5 (KHTML, like Gecko) Chrome/19.0.1084.46 Safari/536.5" 159 host = www2 ; source = tutorialdata.zip;./www2/access.log ; sourcetype = access_combined_wcookie
>	4/28/14 6:20:56.000 PM	182.236.164.11 - - [28/Apr/2014:18:20:56] "GET /cart.do?action=addtocart&itemId=EST-15&productid=85-AG-G09&JSESSIONID=SD6SL8FF10ADFF53101 HTTP 1.1" 200 2252 "http://www.buttercupgames.com/oldlink?itemId=EST-15" Mozilla/5.0 (Macintosh; Intel Mac OS X 10_7_4) AppleWebKit/536.5 (KHTML, like Gecko) Chrome/19.0.1084.46 Safari/536.5" 506 host = www1 source = tutorialdata.zip;./www1/access.log sourcetype = access_combined_wcookie
>	4/28/14 6:20:55.000 PM	182.236.164.11 - - [28/Apr/2014:18:20:56] "POST /oldlink?itemId=EST-18&JSESSIONID=SD6SL8FF10ADFF53101 HTTP 1.1" 408 893 "http://www.buttercupgames.com/product.screen?productId=SF-BVS-G01" Mozilla/5.0 (Macintosh; Intel Mac OS X 10_7_4) AppleWebKit/536.5 (KHTML, like Gecko) Chrome/19.0.1084.46 Safari/536.5" 134 host = www1 source = tutorialdata.zip;./www1/access.log sourcetype = access_combined_wcookie

Fields Sidebar 606

Hide Fields ≡ All Fields

Selected Fields

- a host 3
- a source 3
- a sourcetype 1

Interesting Fields

- a action 5
- # bytes 100+
- a categoryId 8
- a clientip 100+
- # date\_hour 24
- # date\_mday 8
- # date\_minute 60

FIG. 6A

Data Summary				X
Hosts (5)		Sources (8)	Sourcetypes (3)	
<input type="text" value="filter"/>				
Host	Count	Last Update		
mailsv	9,829	4/29/14 1:32:47.000 PM		
vendor_sales	30,244	4/29/14 1:32:46.000 PM		
www1	24,221	4/29/14 1:32:44.000 PM		
www2	22,595	4/29/14 1:32:47.000 PM		
www3	22,975	4/29/14 1:32:45.000 PM		

FIG. 6B






FIG. 7A




FIG. 7B




FIG. 7C


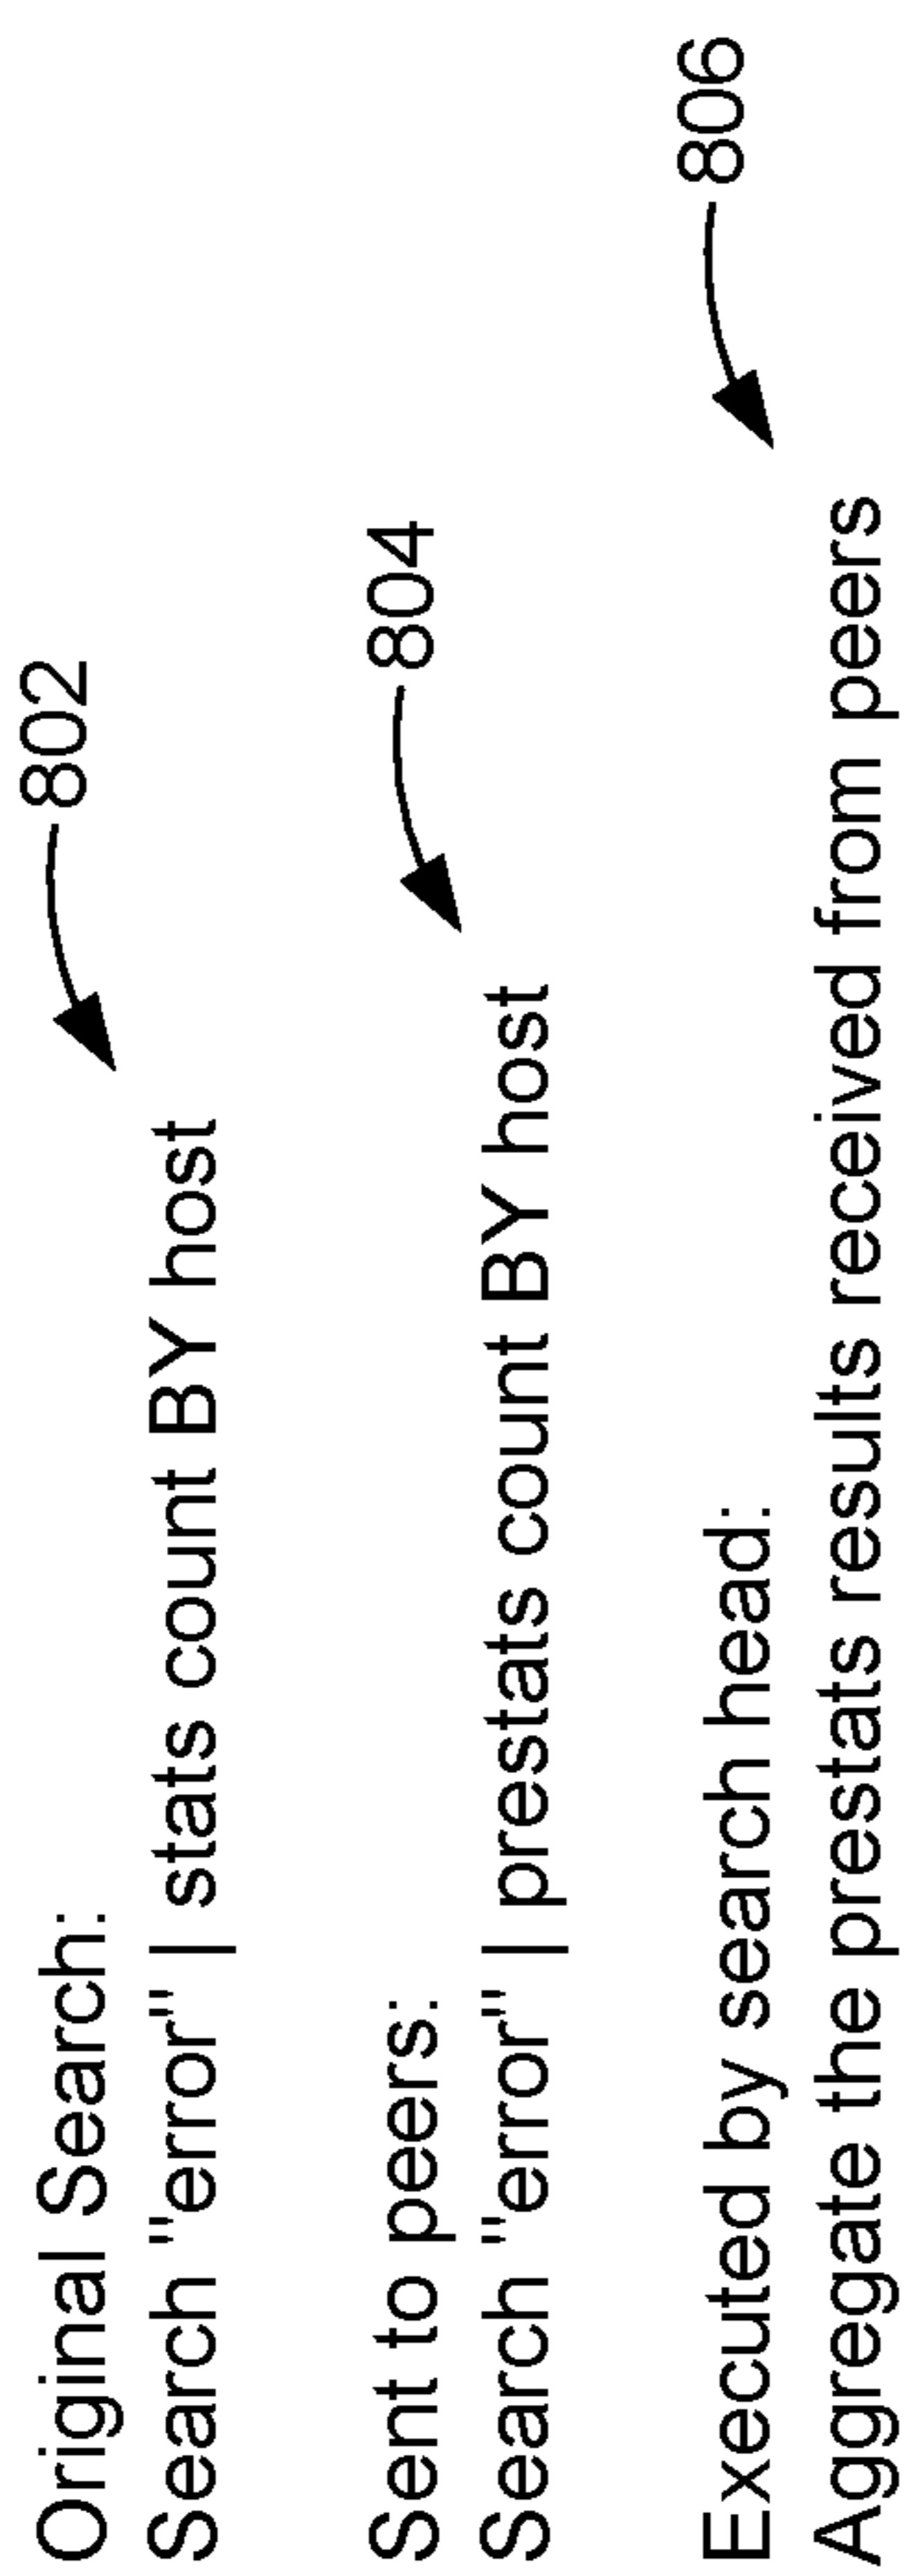


FIG. 7D



**FIG. 8**

KEY INDICATORS VIEW 900


ACCESS NOTABLES Total Count	7 +45	45	7 +15	15	7 +61	61	7 +2	2	7 +32	32
--------------------------------	----------	----	----------	----	----------	----	---------	---	----------	----

901

MALWARE INFECTIONS Total Count	7 +63	632	VULNERABILITIES / HOST AVG Medium Severity Or Higher	1.6 -0.2	1.6	HOSTS FULLY PATCHED Percent Of Total Hosts	78.3% +0.2	78.3%
-----------------------------------	----------	-----	---	-------------	-----	---	---------------	-------

902

903



904




FIG. 9A

INCIDENT REVIEW DATSHBOARD 910

Incident Review | Actions ▾

Status:

Urgency:

Owner:

Security domain:

Governance:

Search:

Title:

24 hour window ▾

Last 15 minutes

Last 60 minutes

Last 4 hours

Last 24 hours

Last 7 days

Last 30 days

Last year

Real-time

Other

All time

Custom time...

**TIME RANGE FIELD 912**

✓ 225 matching events

Hide 🔍 Zoom out 🔍 Zoom to selection [u] Deselect

120 60

4:00 AM Sun Aug 26 2012

6:00 AM 8:00 AM

**225 events in a 24 hour window (real-time) (from 11:29:20 AM August 25 to 11:29:20 AM August 26, 2012)**

Select all | Unselect all 1 2 3 4 5 6 7 8 9 10 next | Edit selected events | Edit all 225 matching

INCIDENT ATTRIBUTE FIELDS 911

Linear scale ▾ 1 bar = 1 hour


TIMELINE 913

120 60

EVENTS LIST 914

Select	Options	Time	Security Domain	Title	Urgency	Status	Owner
<input type="checkbox"/>	▾	8/26/12 11:11:03.000 AM	Access ▾	Insecure Or Cleartext Authentication Detected ▾	High ▾	New ▾	unassigned ▾
<input type="checkbox"/>	▾	8/26/12 11:10:07.000 AM	Access ▾	Insecure Or Cleartext Authentication Detected ▾	High ▾	New ▾	unassigned ▾
<input type="checkbox"/>	▾	8/26/12 11:00:39.000 AM	Access ▾	Account (blinebry) Deleted On (PROD-POS-001) ▾	High ▾	New ▾	unassigned ▾
<input type="checkbox"/>	▾	8/26/12 11:00:39.000 AM	Access ▾	Account (beu) Deleted On (COREDEV-006) ▾	High ▾	New ▾	unassigned ▾
<input type="checkbox"/>	▾	8/26/12 11:00:39.000 AM	Access ▾	Account (combs) Deleted On (HOST-005) ▾	High ▾	New ▾	unassigned ▾
<input type="checkbox"/>	▾	8/28/12	Access ▾	Account (wisner) Deleted On (BUSDEV-005) ▾	High ▾	New ▾	unassigned ▾

FIG. 9B



**FIG. 9C**




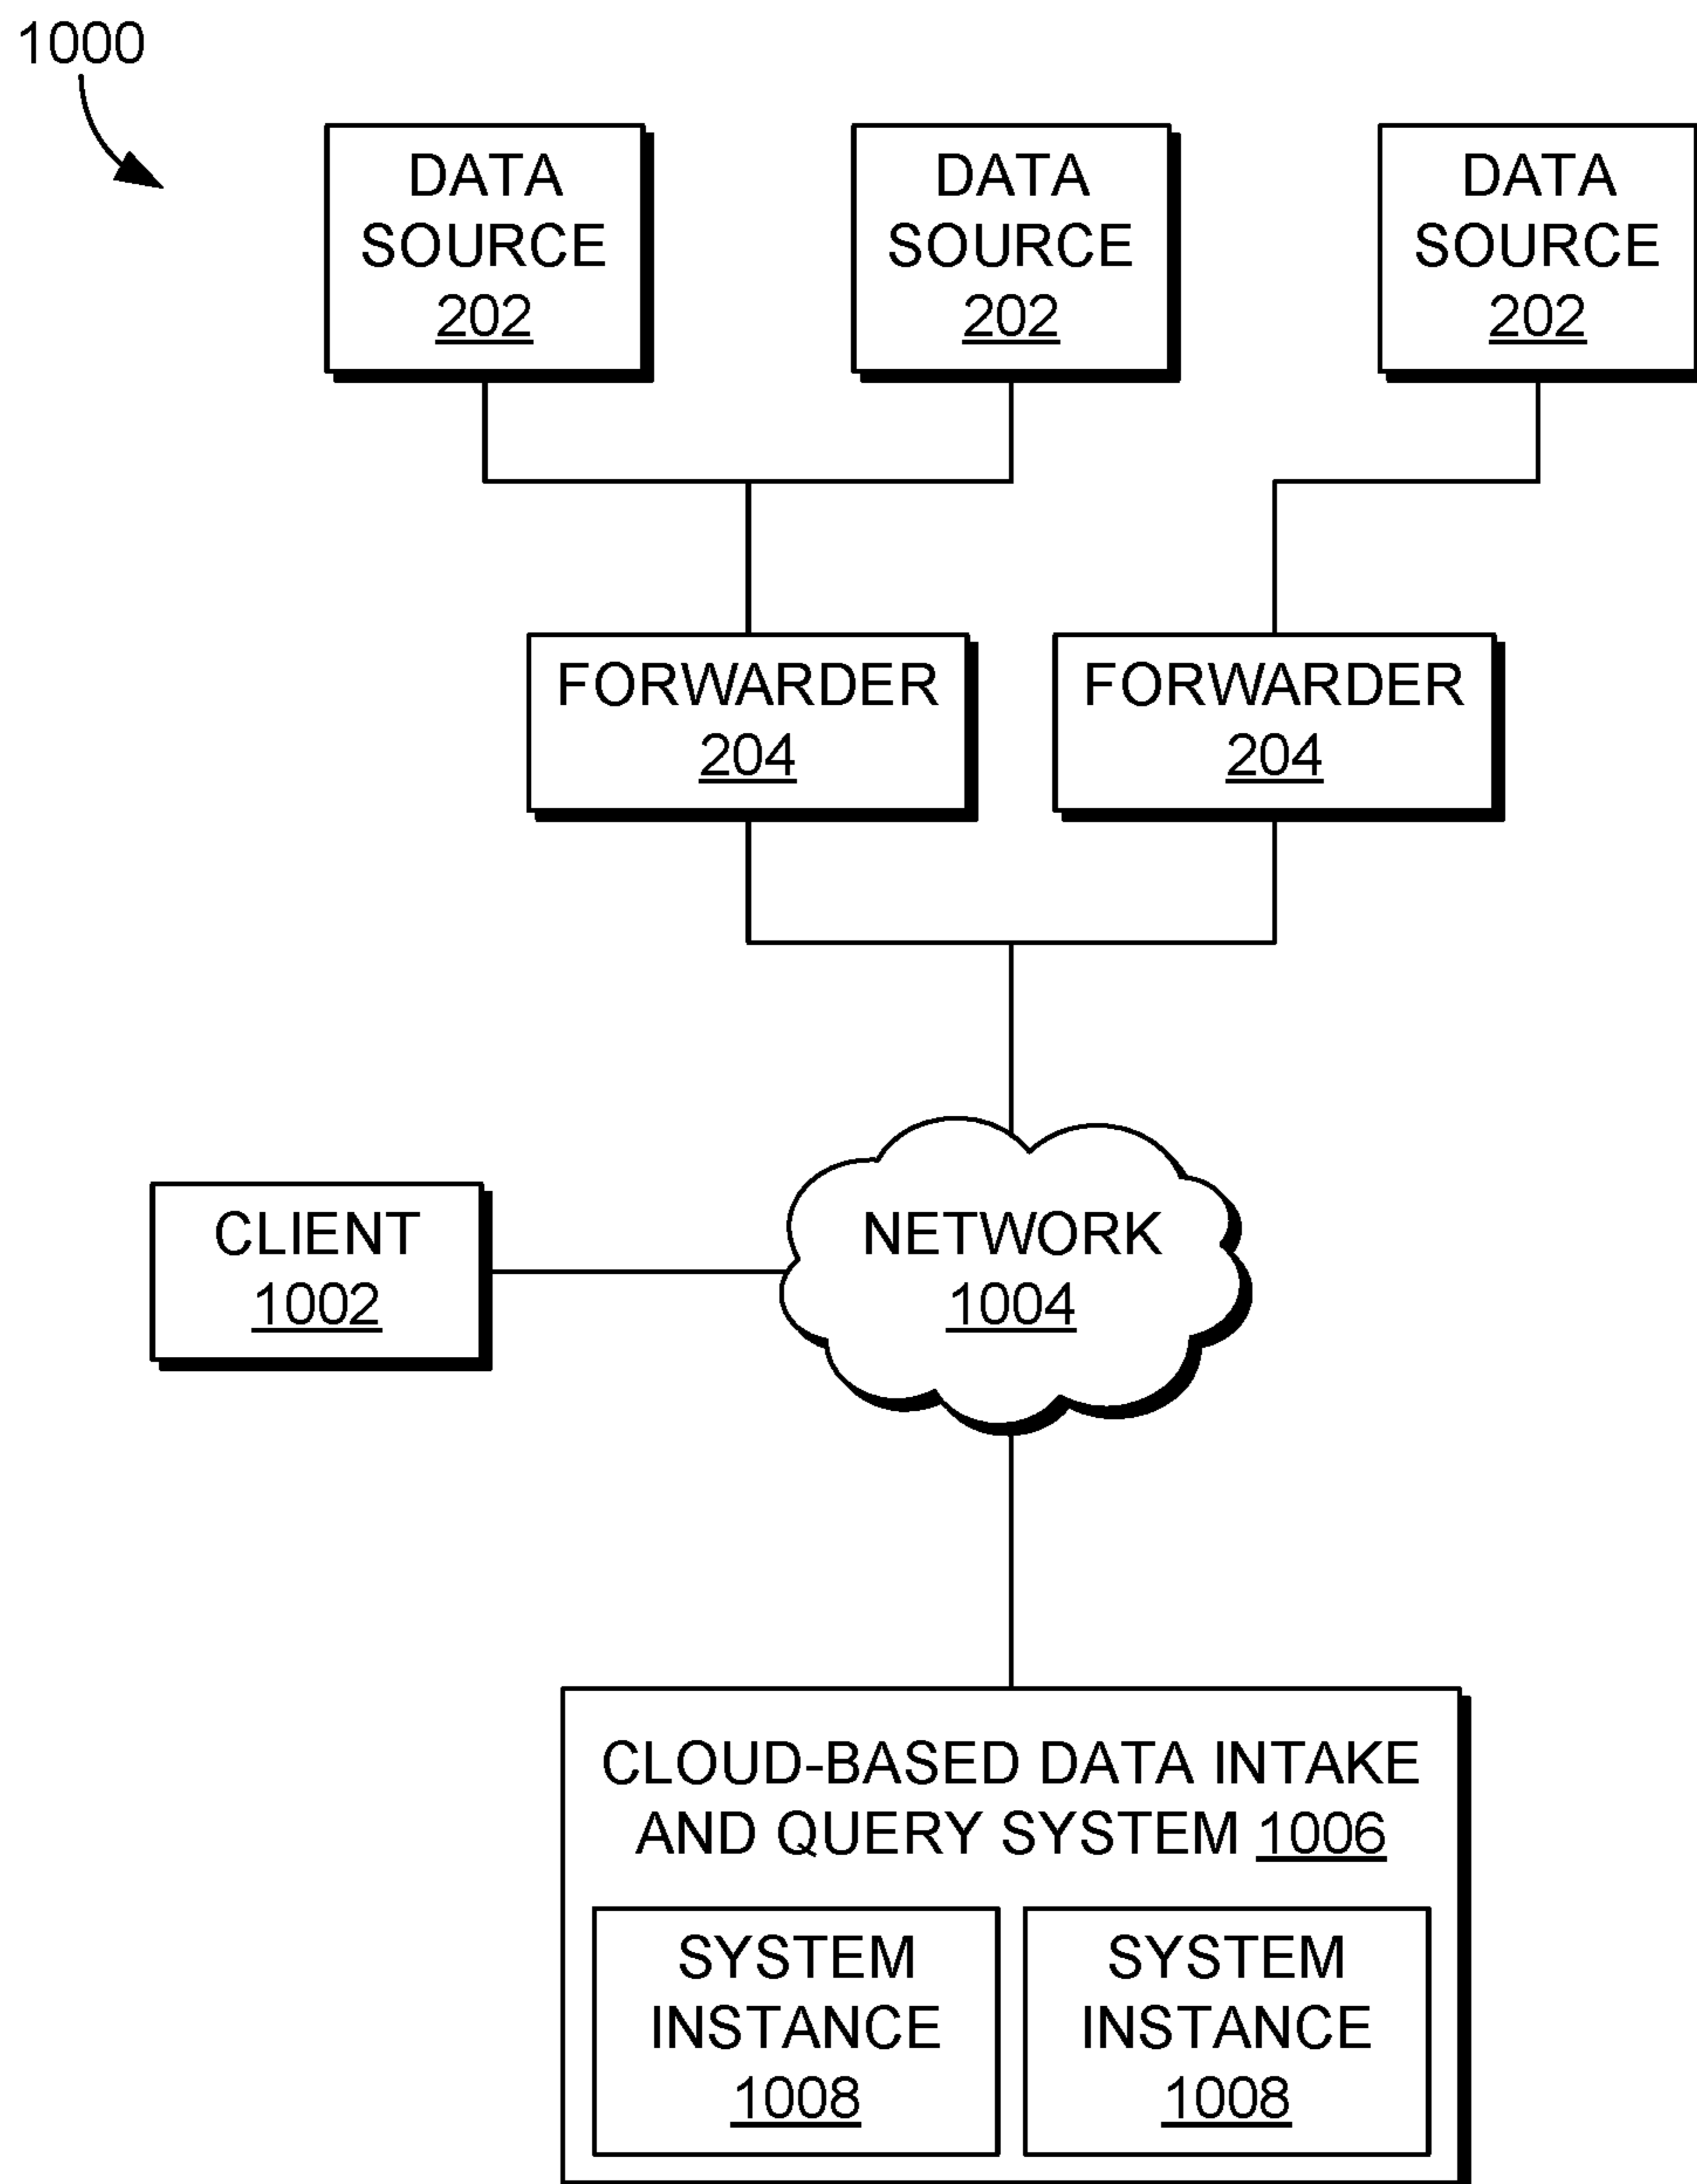


FIG. 9D



**FIG. 10**

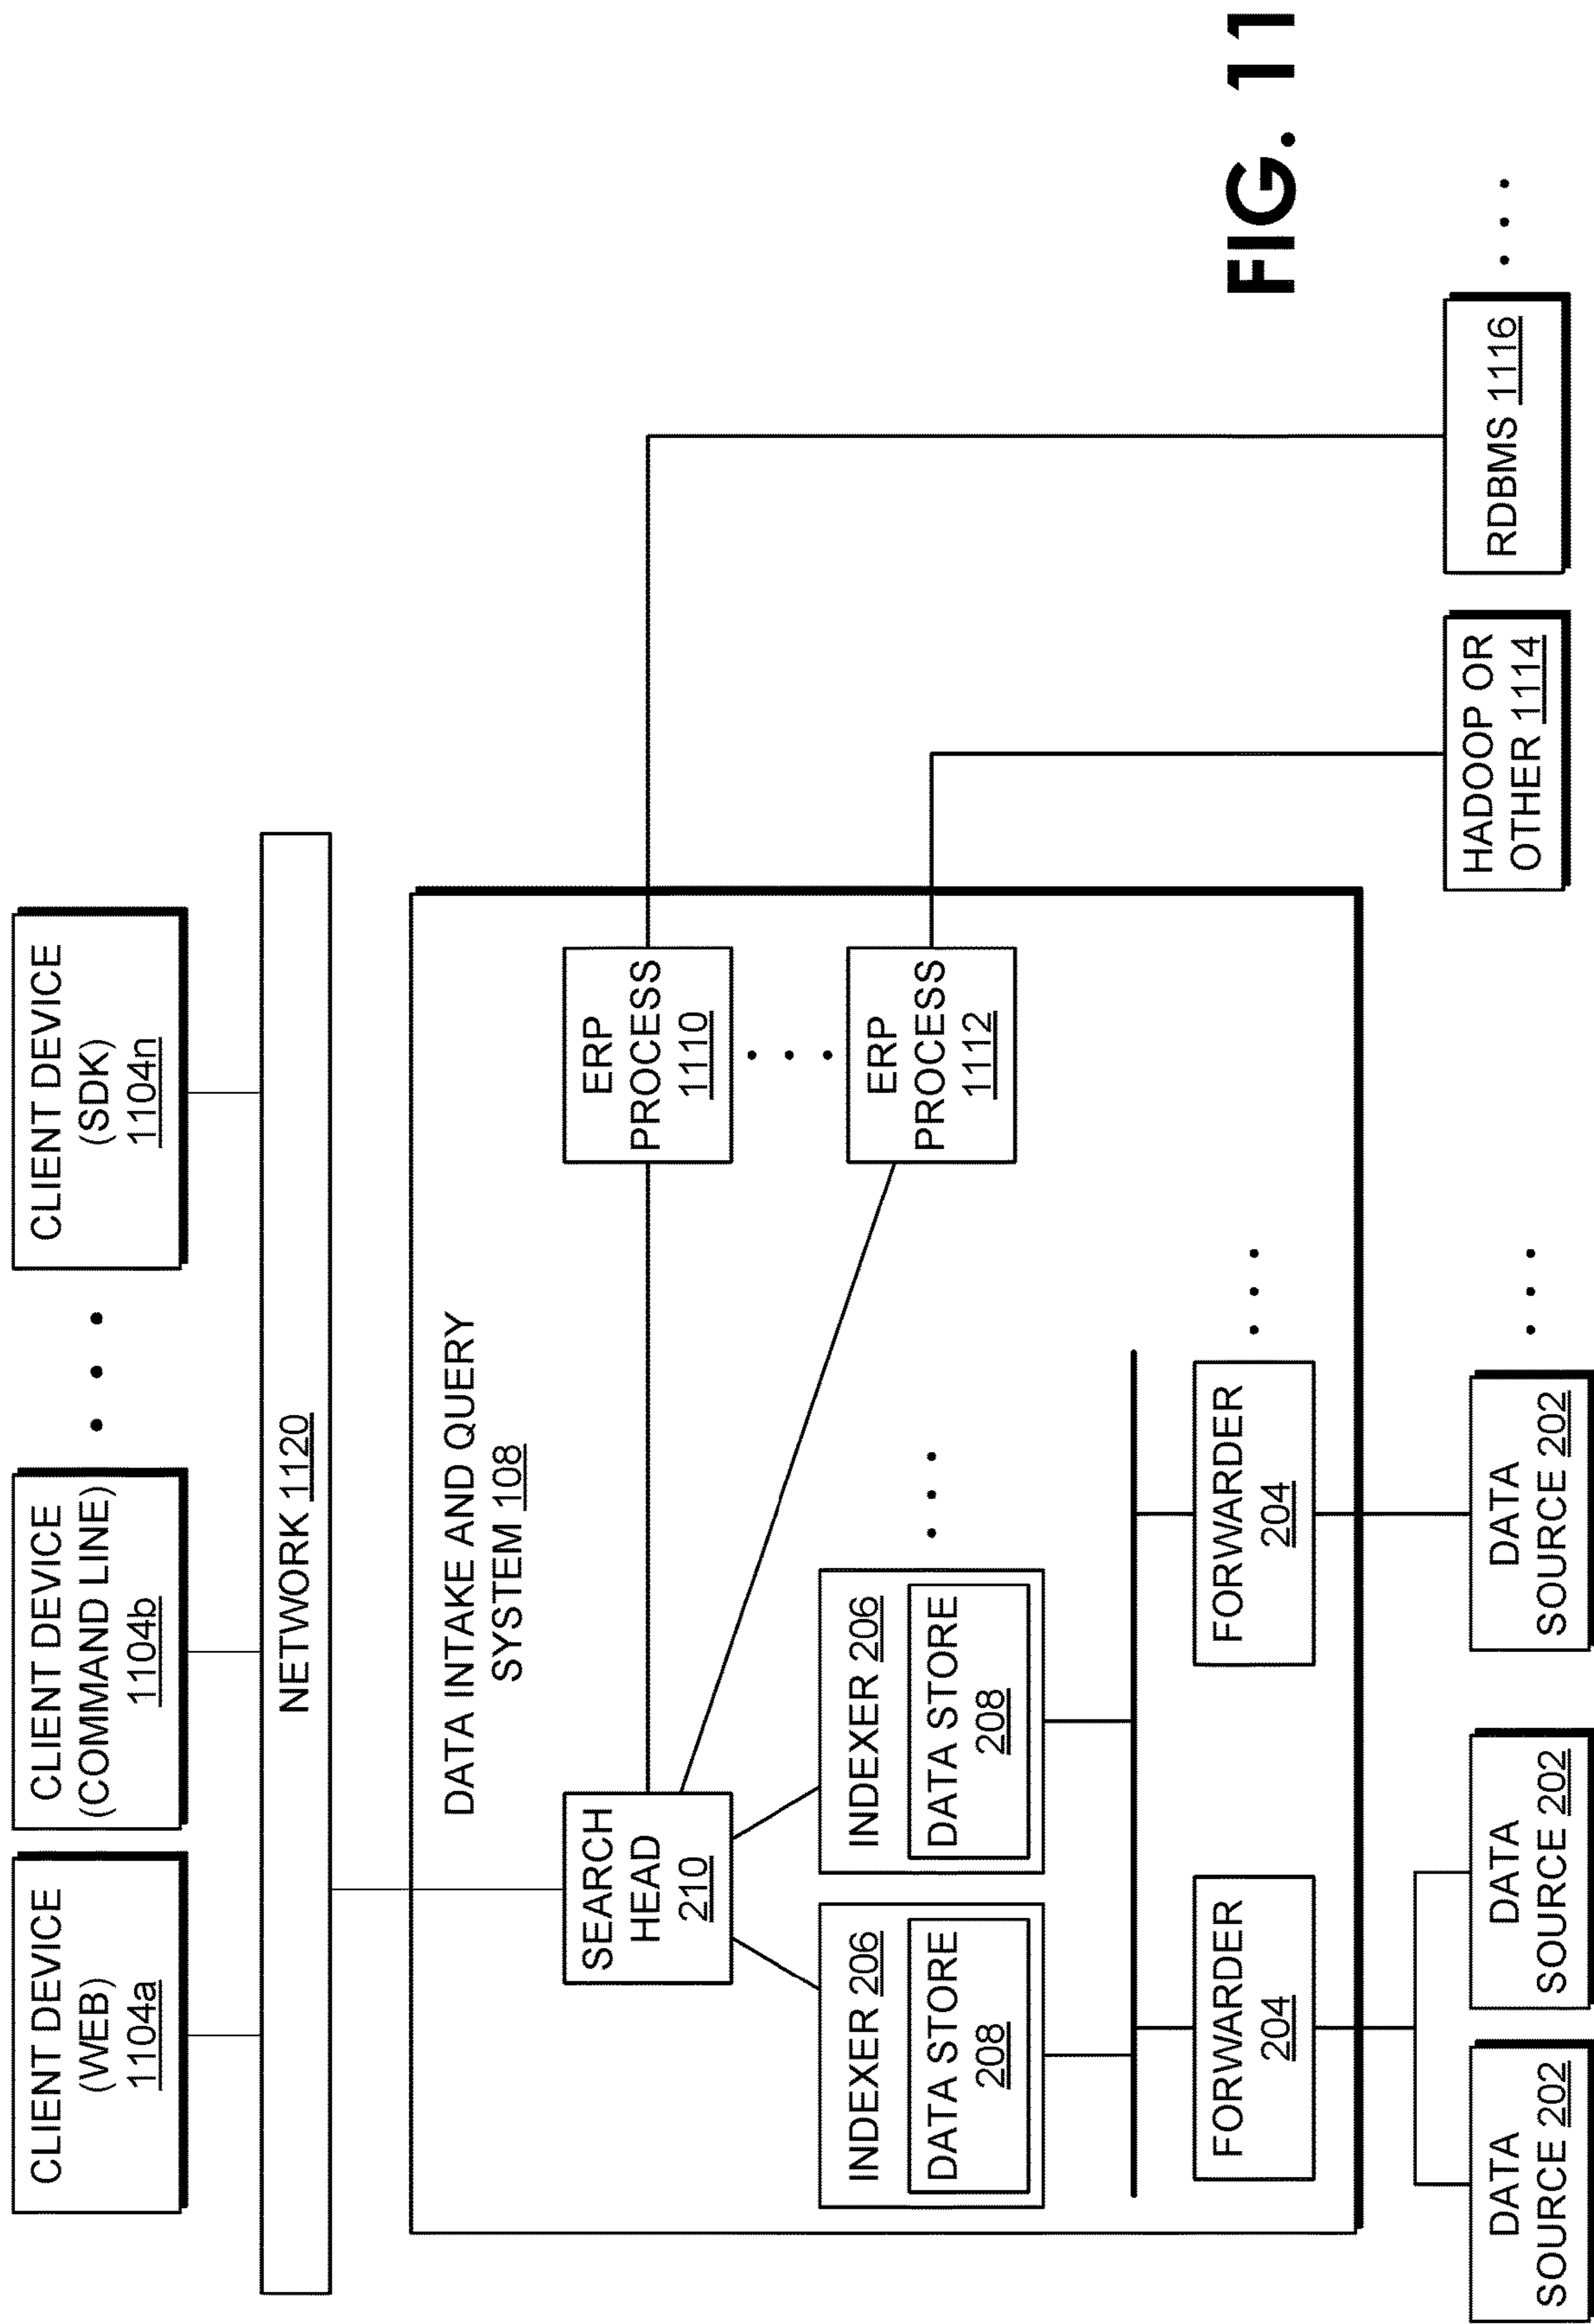


FIG. 11

1200

Select a Data Model	
<i>i</i>	4 Data Models ~1201
▶	<u>Buttercup Games Sales</u> ~1202
▶	Splunk's Internal Audit Logs - SAMPLE
▶	Splunk's Internal Server Logs - SAMPLE
▶	test

**FIG. 12**

1300

Select an Object	
◀ Back	
<i>i</i>	6 Objects in Buttercup Game Sales ~1301
▶	Buttercup Games Web Store Events
▶	HTTP Success
▶	<u>Successful Purchases</u> ~1302
▶	Failed Purchases
▶	HTTP Client Error
▶	HTTP Server Error

**FIG. 13**

New Pivot

Save As... ▾ Clear Successful\_purchases ▾

---

1,966 events (before 9/22/13 5:19:02.000 PM)

**Filters**

All time + Highest 10 product name by price

Split Rows + Highest 10 produ...

product name + price

**Split Columns** + Documentation

Column Values + Sum of price

Count of Successful purchases + Sum of price

---

product name	price	Count of Successful purchases	Sum of price
Dream Crusher	39.99	227	9077.73
Manganiello Bros.	39.99	199	7958.01
Orvil the Wolverine	39.99	169	6758.31
World of Cheese	24.99	267	6672.33
Mediocre Kingdoms	24.99	250	6247.50
SIM Cubicle	19.99	233	4657.67
Final Sequel	24.99	181	4523.19
Benign Space Debris	24.99	135	3373.65
Curling 2014	19.99	126	2518.74
Holy Blade of Gouda	5.99	179	1072.21

FIG. 14




FIG. 15




FIG. 16




FIG. 17



**Search**

source`type=access_combined`  
| `f`

**1806A**

<code>table time, place, lat*, lon*, mag</code>	Recently Used
<code>timechart per_hour(price) by product_name usenull=f</code>	Recently Used
<code>top limit=20 referer</code>	Recently Used
<code>table &lt;wc-field-list&gt;</code>	Command
<code>tags [outputfield=&lt;field&gt;] [incln...value=&lt;bool&gt;] &lt;field-list&gt;</code>	Command
<code>tail [&lt;N&gt;]</code>	Command
<code>timechart [sep=&lt;string&gt;] [form...n&gt;) BY &lt;split-by-clause&gt; )</code>	Command
<code>top [&lt;n&gt;] [&lt;top-options&gt;...] &lt;field-list&gt; [BY &lt;by-clause&gt;]</code>	Command

**What to Search**

246,225 Events INDEXED      23 days ago EARLIEST EVENT      a minute ago LATEST EVENT

**1802**

All Time ▾      Verbose Mode ▾

FIG. 18A

**Search**

source`type=access_combined`  
| `f`

**1806B**

<code>timechart per_hour(price) by product_name usenull=f</code>	Recently Used
<code>timechart [sep=&lt;string&gt;] [form...n&gt;) BY &lt;split-by-clause&gt; )</code>	Command

Creates a time series chart with corresponding table of statistics. Show More

**1802**

All Time ▾      Verbose Mode ▾

FIG. 18B

**Search**

source`type=access_combined`  
| `timechart [sep=<string>] [format=<string>] [partial=<bool>] [cont=<bool>] [limit=<int>] [<agg>] [<bin-options>...] ( (<single-agg> [BY <split-by-clause>] ) | (<eval-ex`

Used to construct output field names when multiple data series are used in conjunctions with a split-by field.

**1810**

All Time ▾      Verbose Mode ▾

FIG. 18C

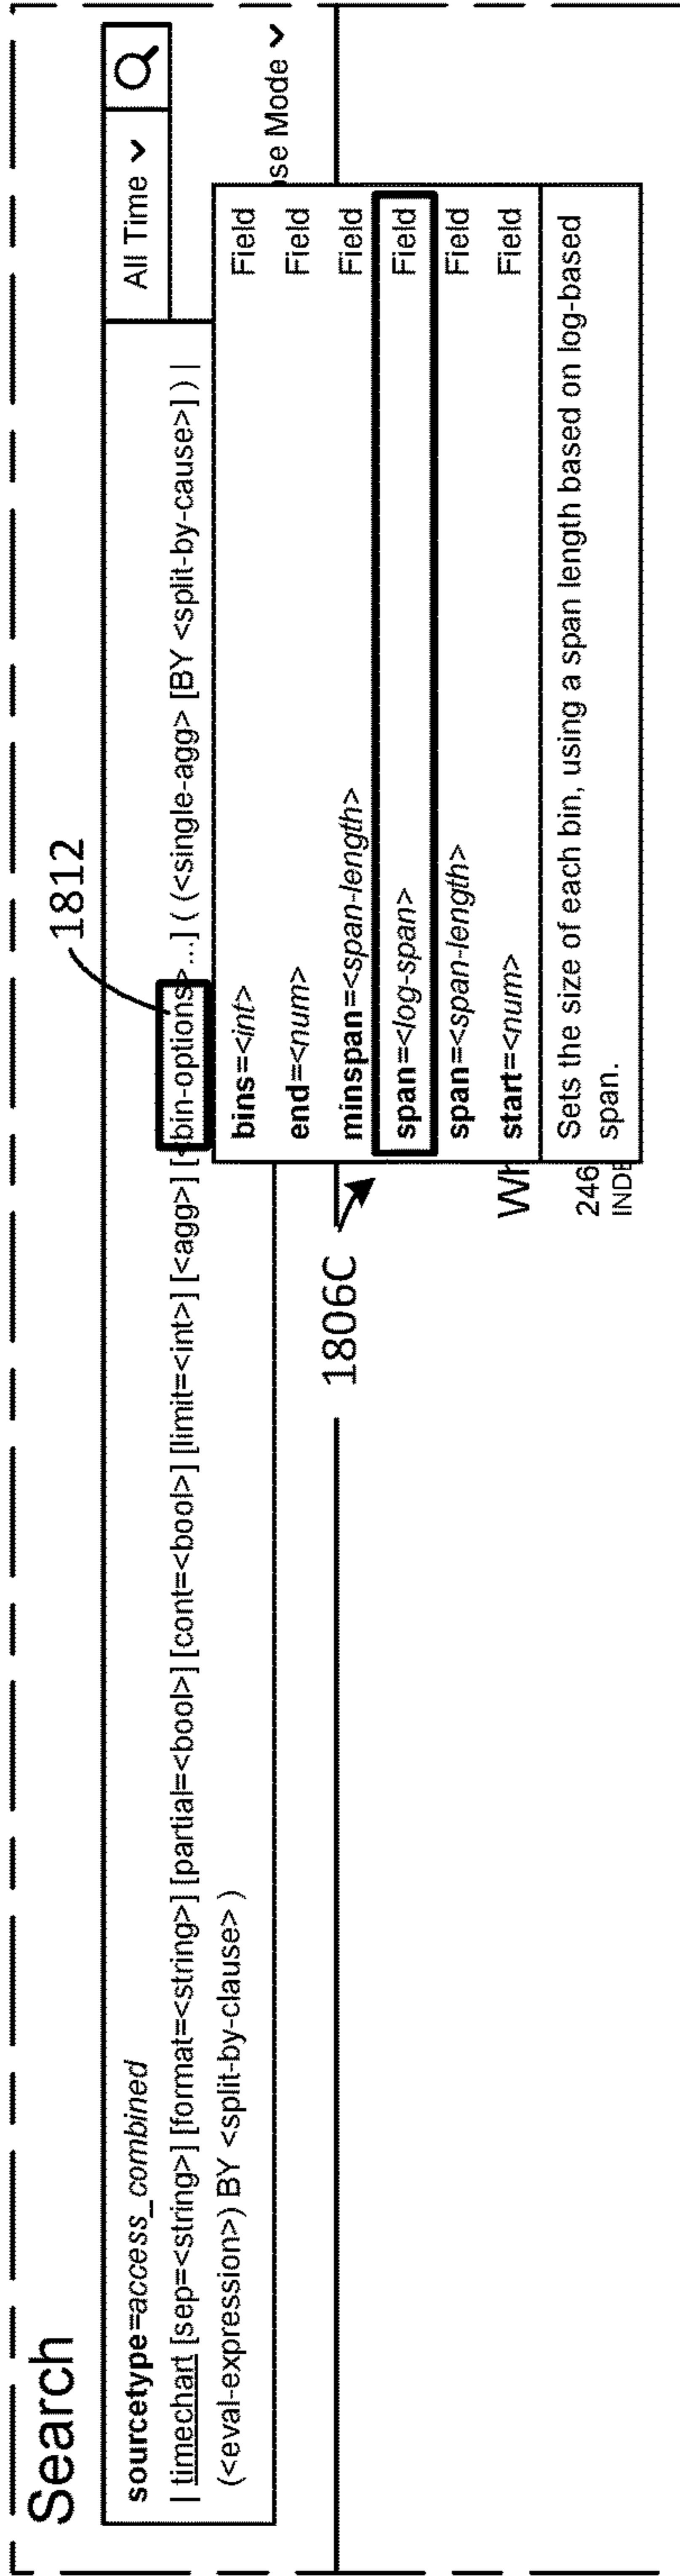


FIG. 18D




FIG. 18E

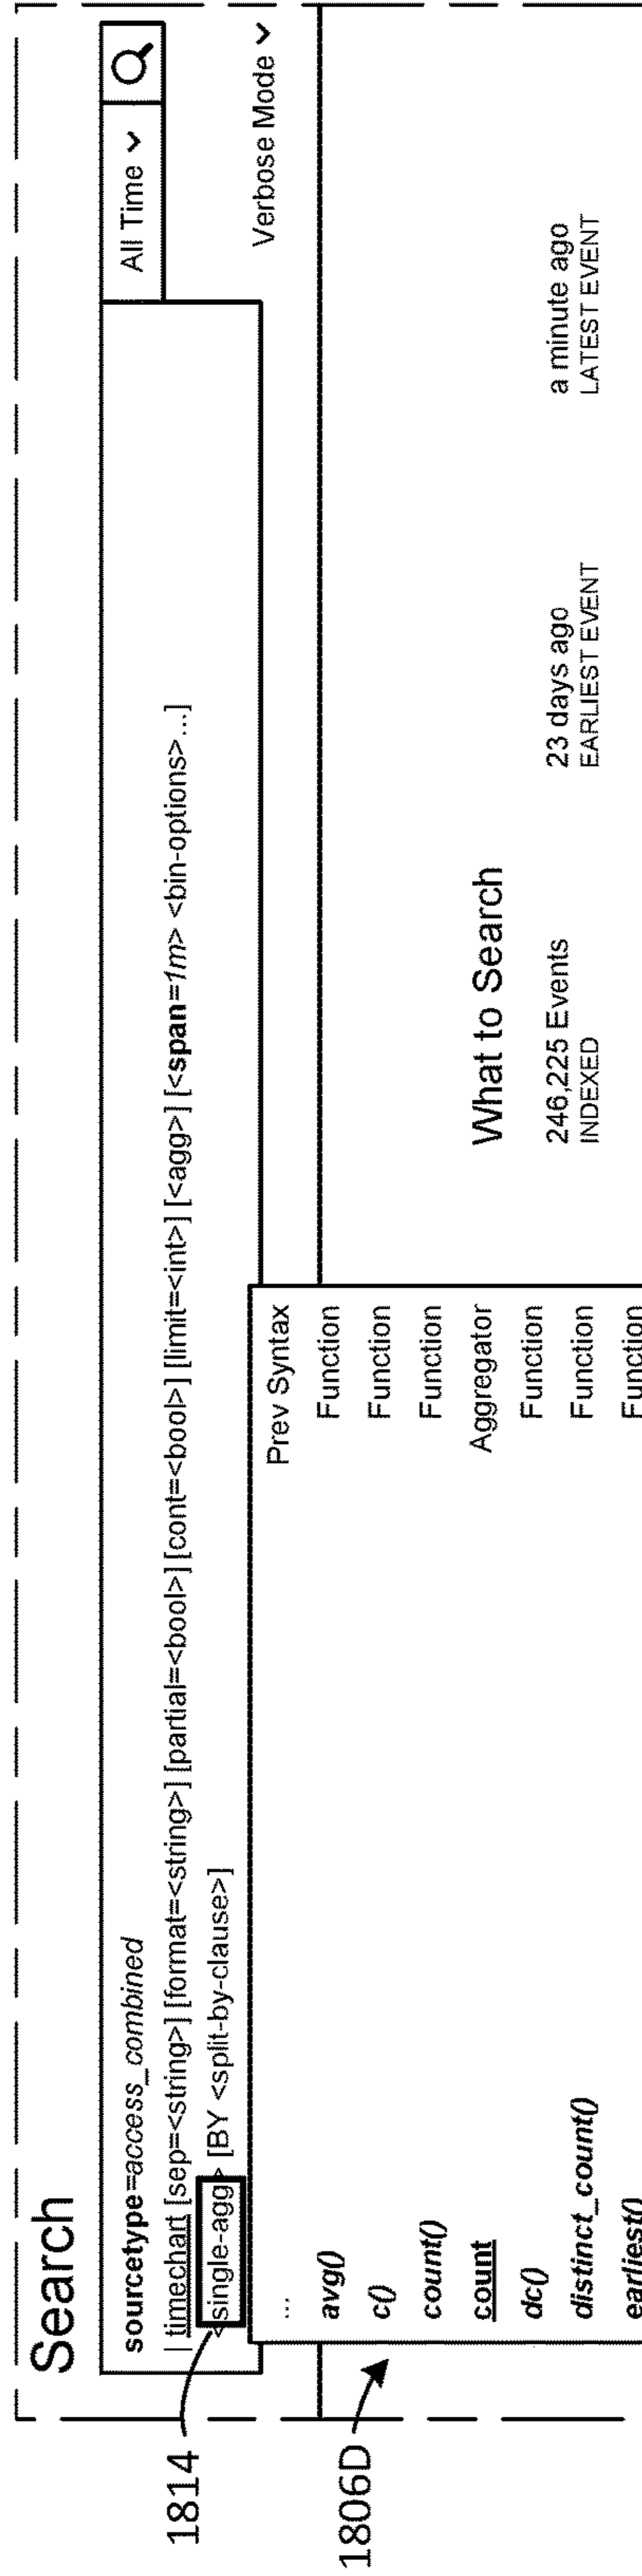


FIG. 18F




FIG. 18G

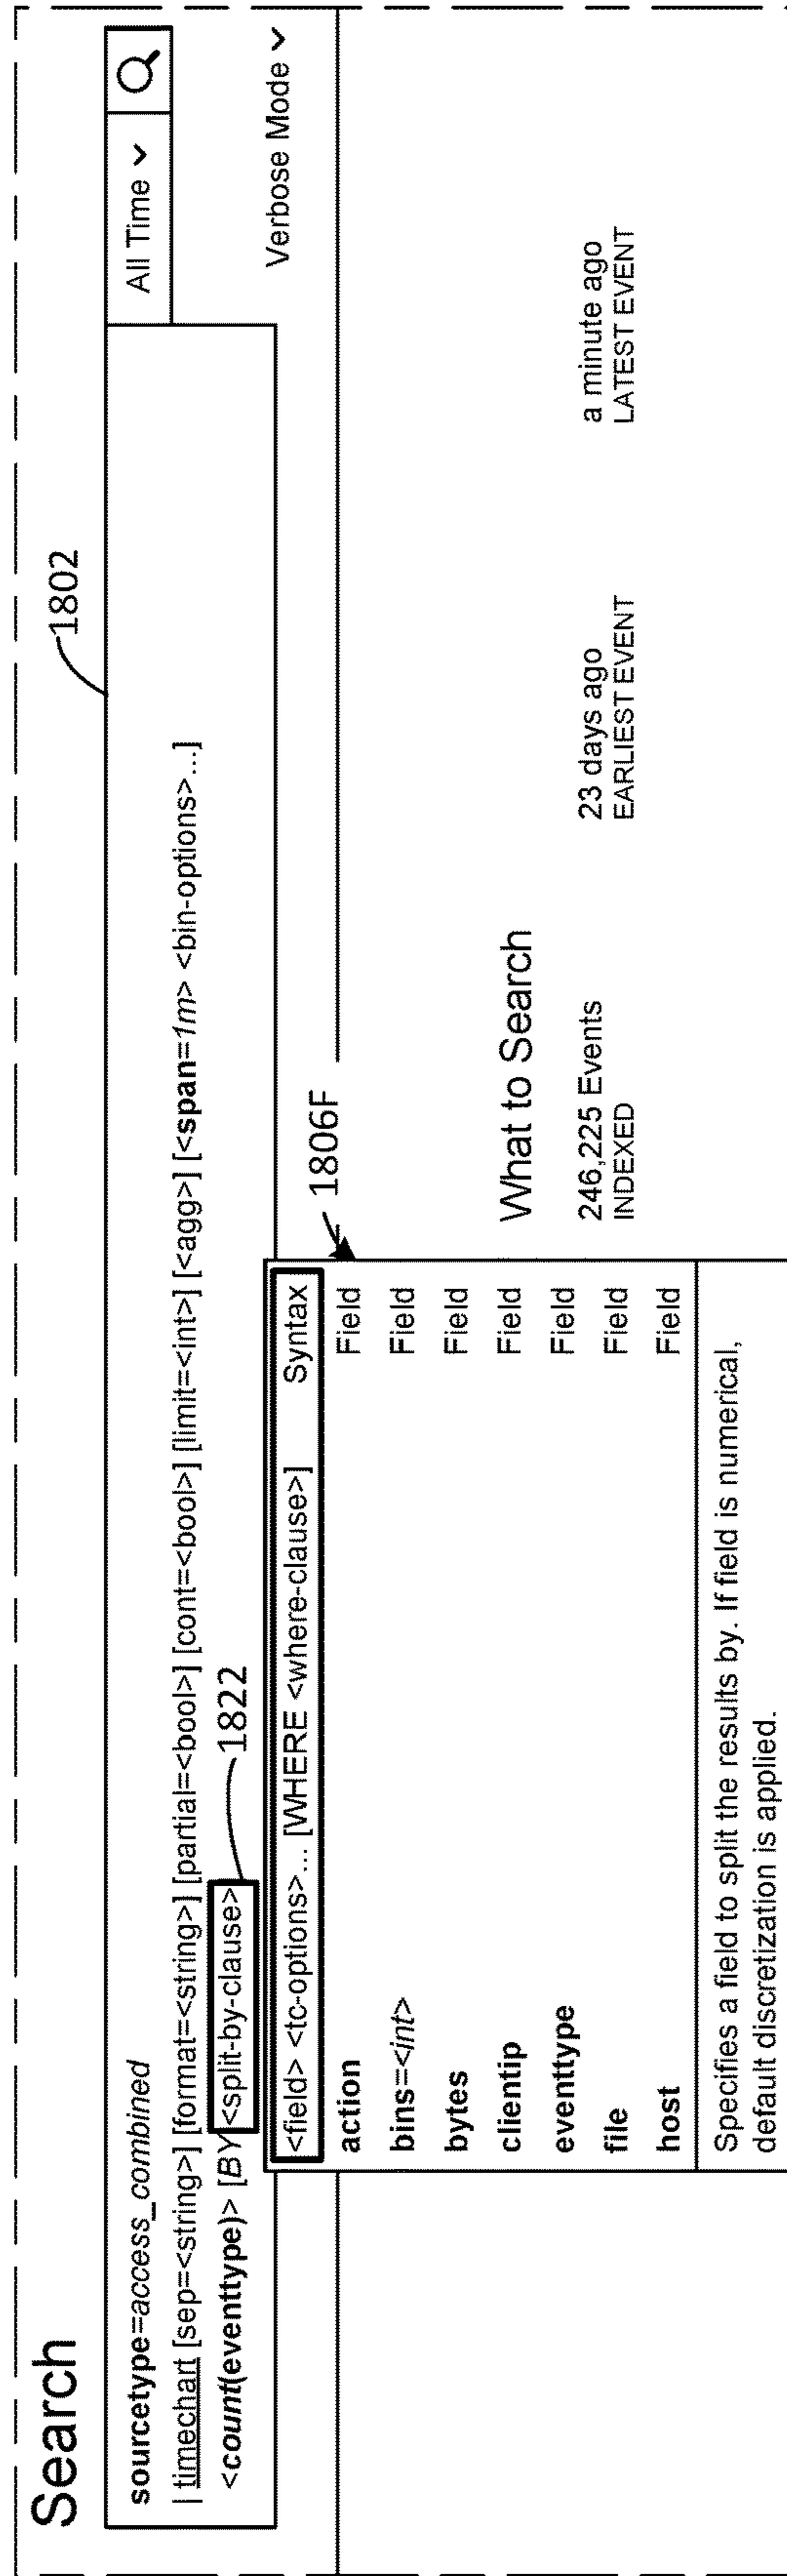


FIG. 18H




FIG. 18I




FIG. 18J




FIG. 18K

**Search**

`sourcetype=access_combined`  
| timechart [sep=<string>] [format=<string>] [partial=<bool>] [cont=<bool>] [limit=<int>] [<agg>] [<span=1m> <bin-options> ...]  
<count(eventtype)> [BY <source\_ip> <usnull=f> <tc-options> ...] [WHERE <count> <wherethresh-comp>]]

All Time

Verbose Mode

FIG. 18L

**Search**

`sourcetype=access_combined`  
| timechart [sep=<string>] [format=<string>] [partial=<bool>] [cont=<bool>] [limit=<int>] [<agg>] [<span=1m> <bin-options> ...]  
<count(eventtype)> [BY <source\_ip> <usnull=f> <tc-options> ...] [WHERE <count> >10]]

All Time

Verbose Mode

FIG. 18M

**Search**

`sourcetype=access_combined`  
| timechart span=1m count(eventtype) BY source\_ip usnull=f WHERE count >10]

All Time

Verbose Mode

FIG. 18N

```

this.$rules = {
  ...
  foo-command : [
    //double quotes
    {
      token: 'text',
      regex: /"(?:\\|.)"?"/
    },
    //pipe
    {
      token: 'pipe',
      regex: /\|/,
      next: 'command'
    },
    //start subsearch
    {
      token: 'subsearch',
      regex: /\[/,
      push: 'command'
    },
    //end subsearch
    {
      token: 'subsearch',
      regex: /\]/,
      next: 'pop'
    },
    //args
    {
      token: ['argument', 'operator'],
      regex: /\b(bar)(=)/,
      push: 'boolState'
    },
    {
      token: ['argument', 'operator'],
      regex: /\b(baz)(=)/,
      push: 'intState'
    },
    //functions
    {
      token: ['function', 'text'],
      regex: /\b(count|avg|abs)(\s*(\s*)/
    },
    //keywords
    {
      token: 'modifier',
      regex: /\b(BY)\b/
    },
    //other
    {the token for ints}
  ],
  ...
}

```


FIG. 19B

```

searchbnf.conf
[command-foo]
syntax = foo (<foo-options>)* <foo-
functions> " "<field> "j" <by-int-phrase>
[foo-options]
syntax = <bax-option> | <bar-option>
[bax-option]
syntax = baz=<int>
[bar-option]
syntax = bar=<bool>
[foo-functions]
syntax = count|avg|abs
[by-int-phrase]
syntax = BY <int>

```

FIG. 19A



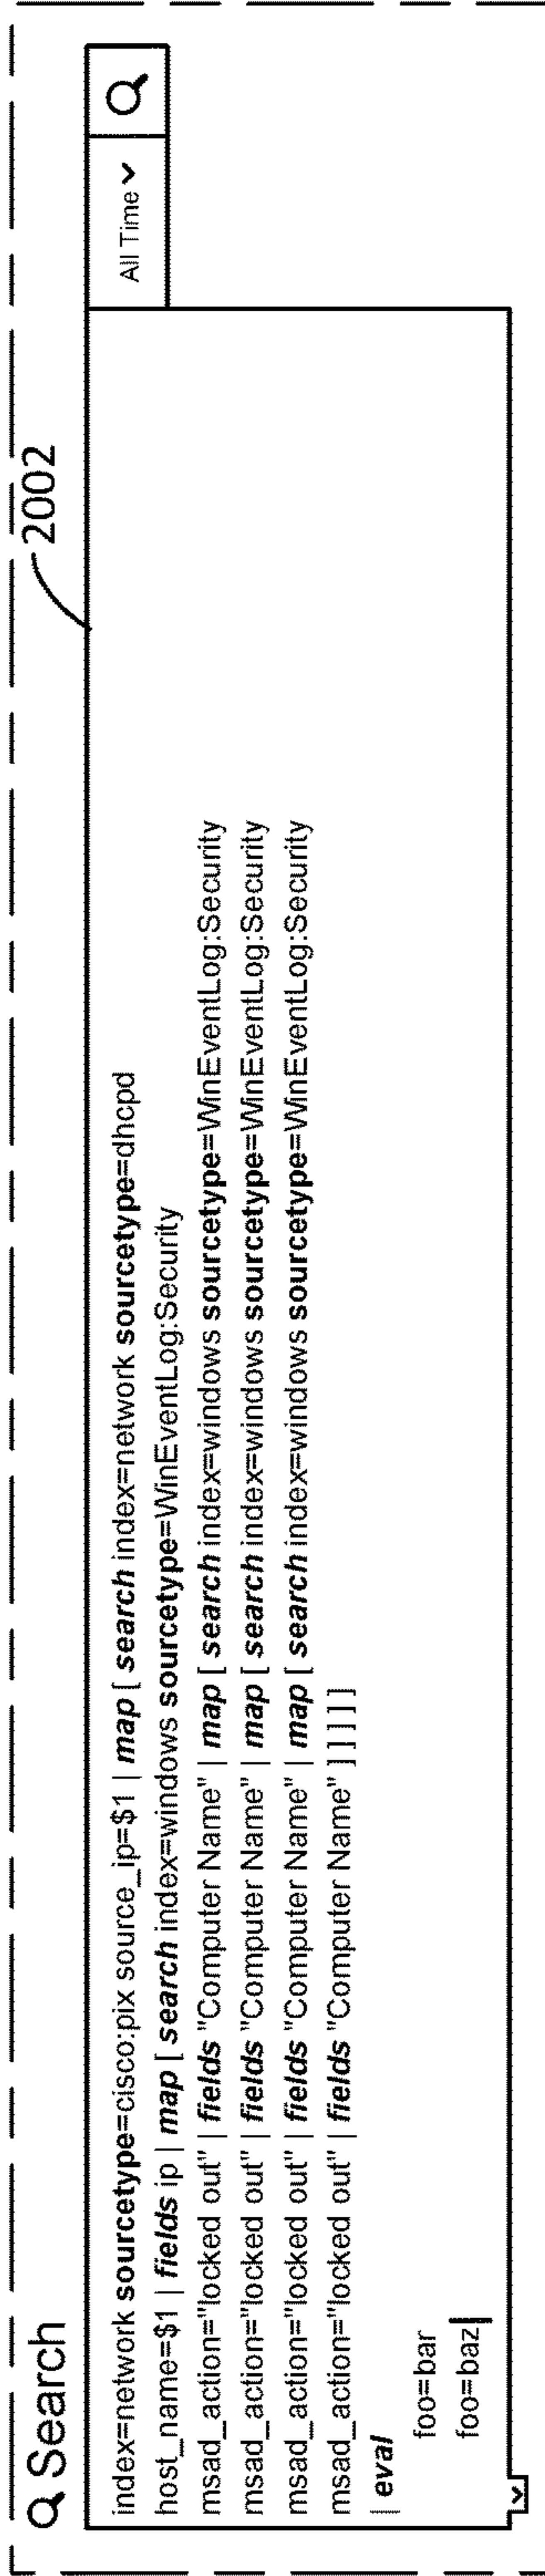


FIG. 20A

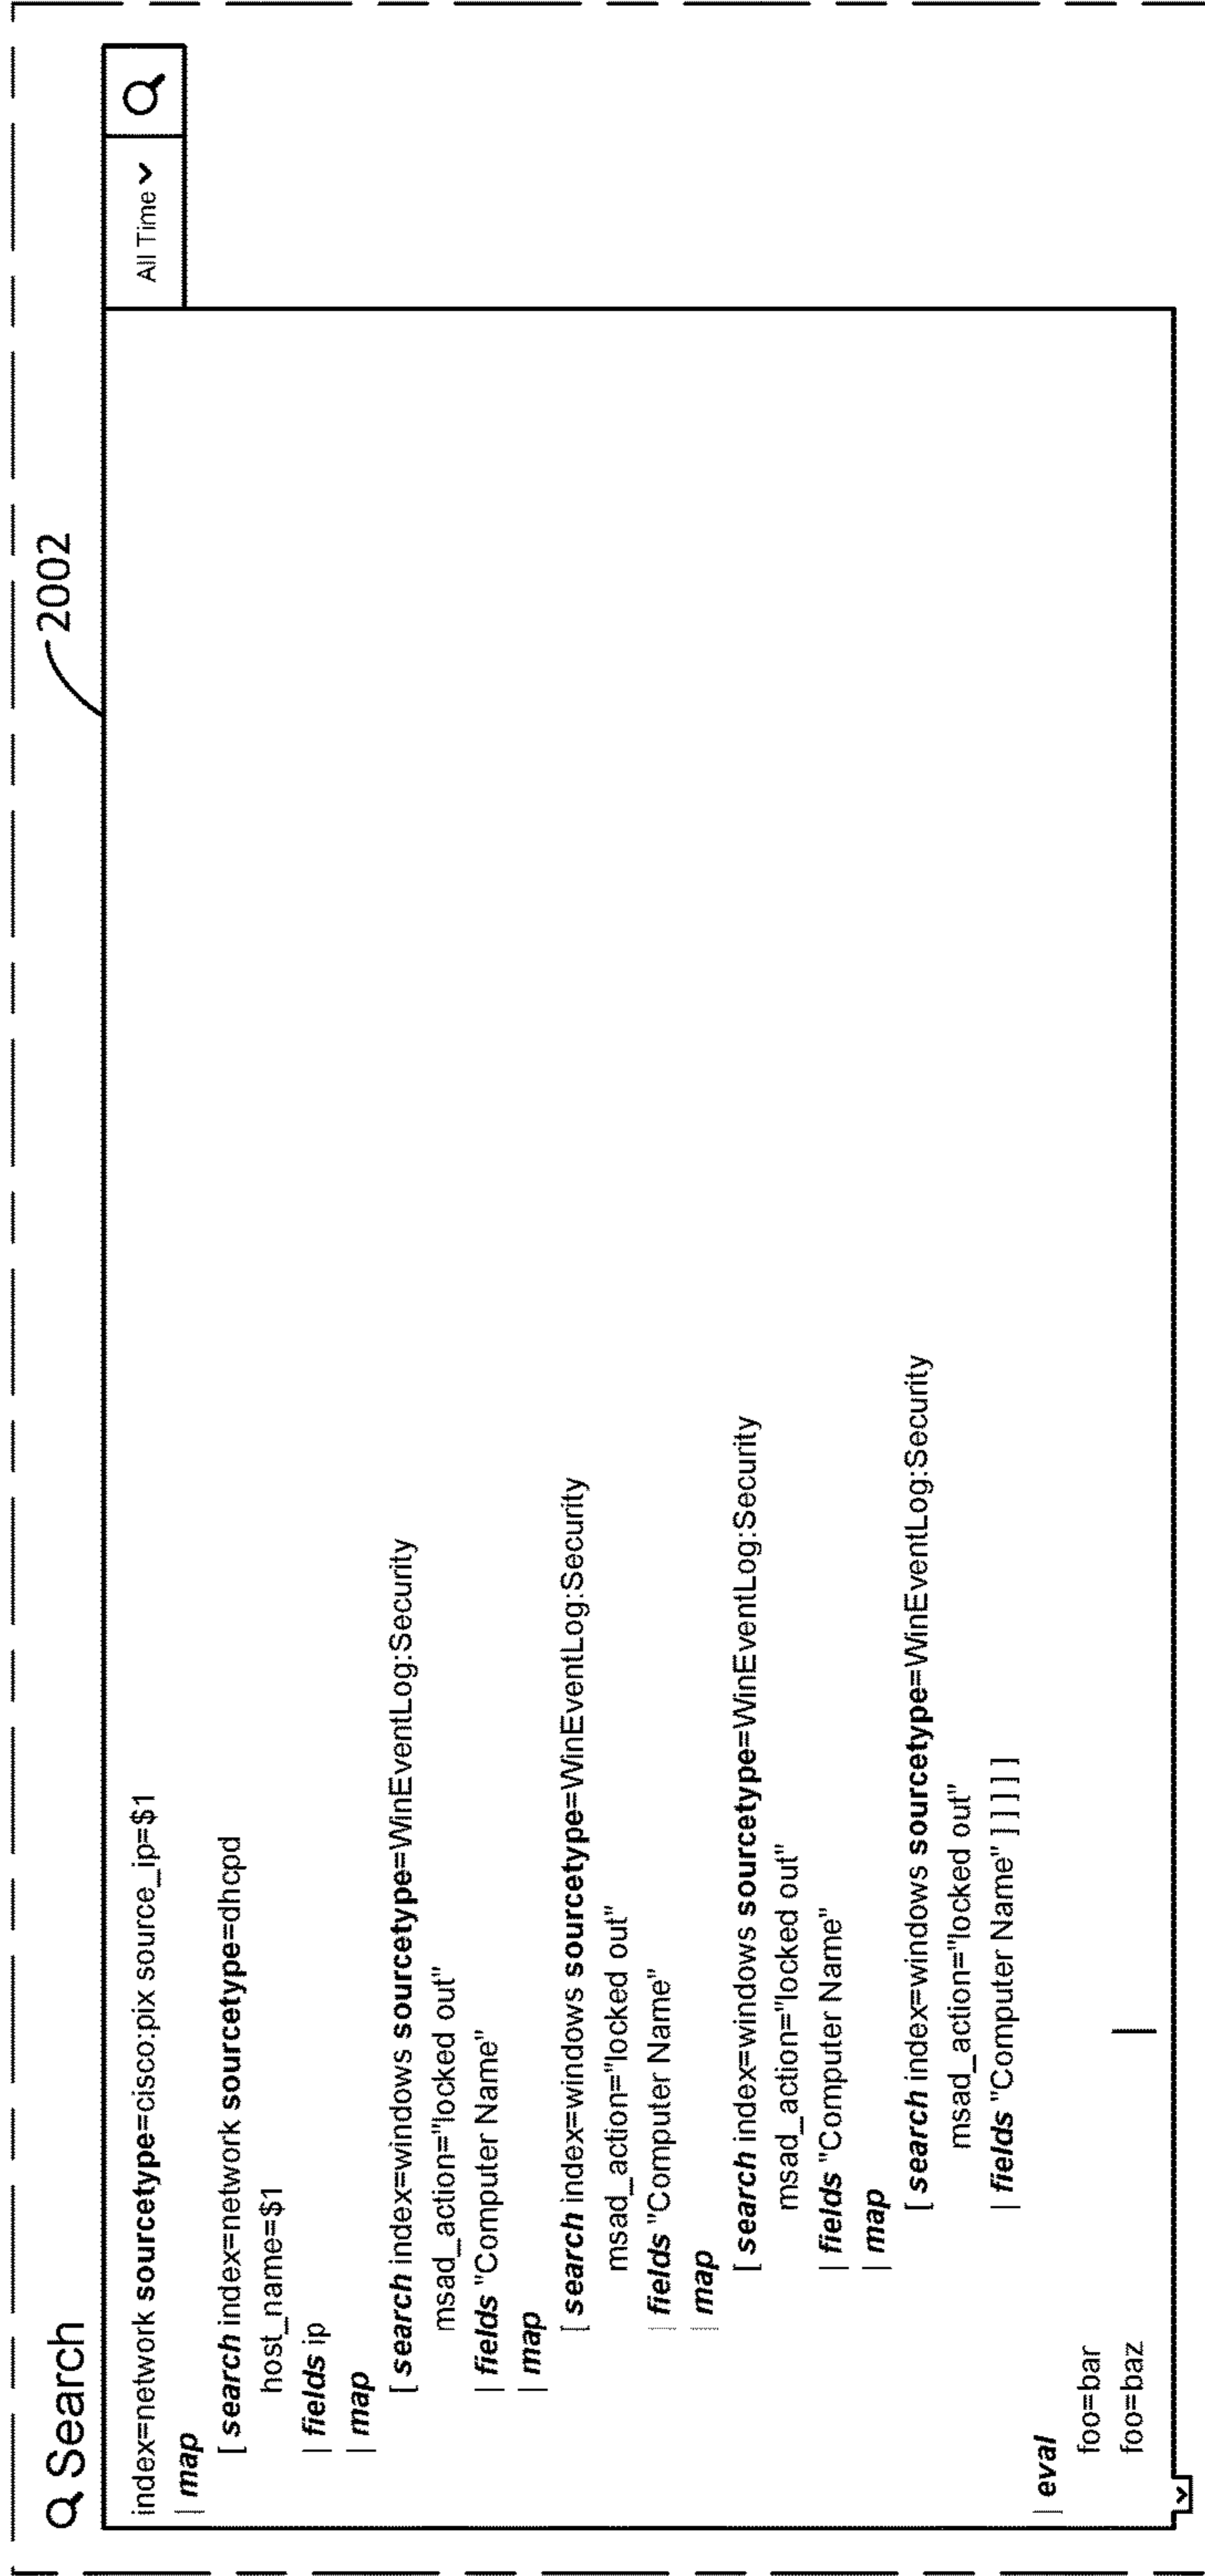

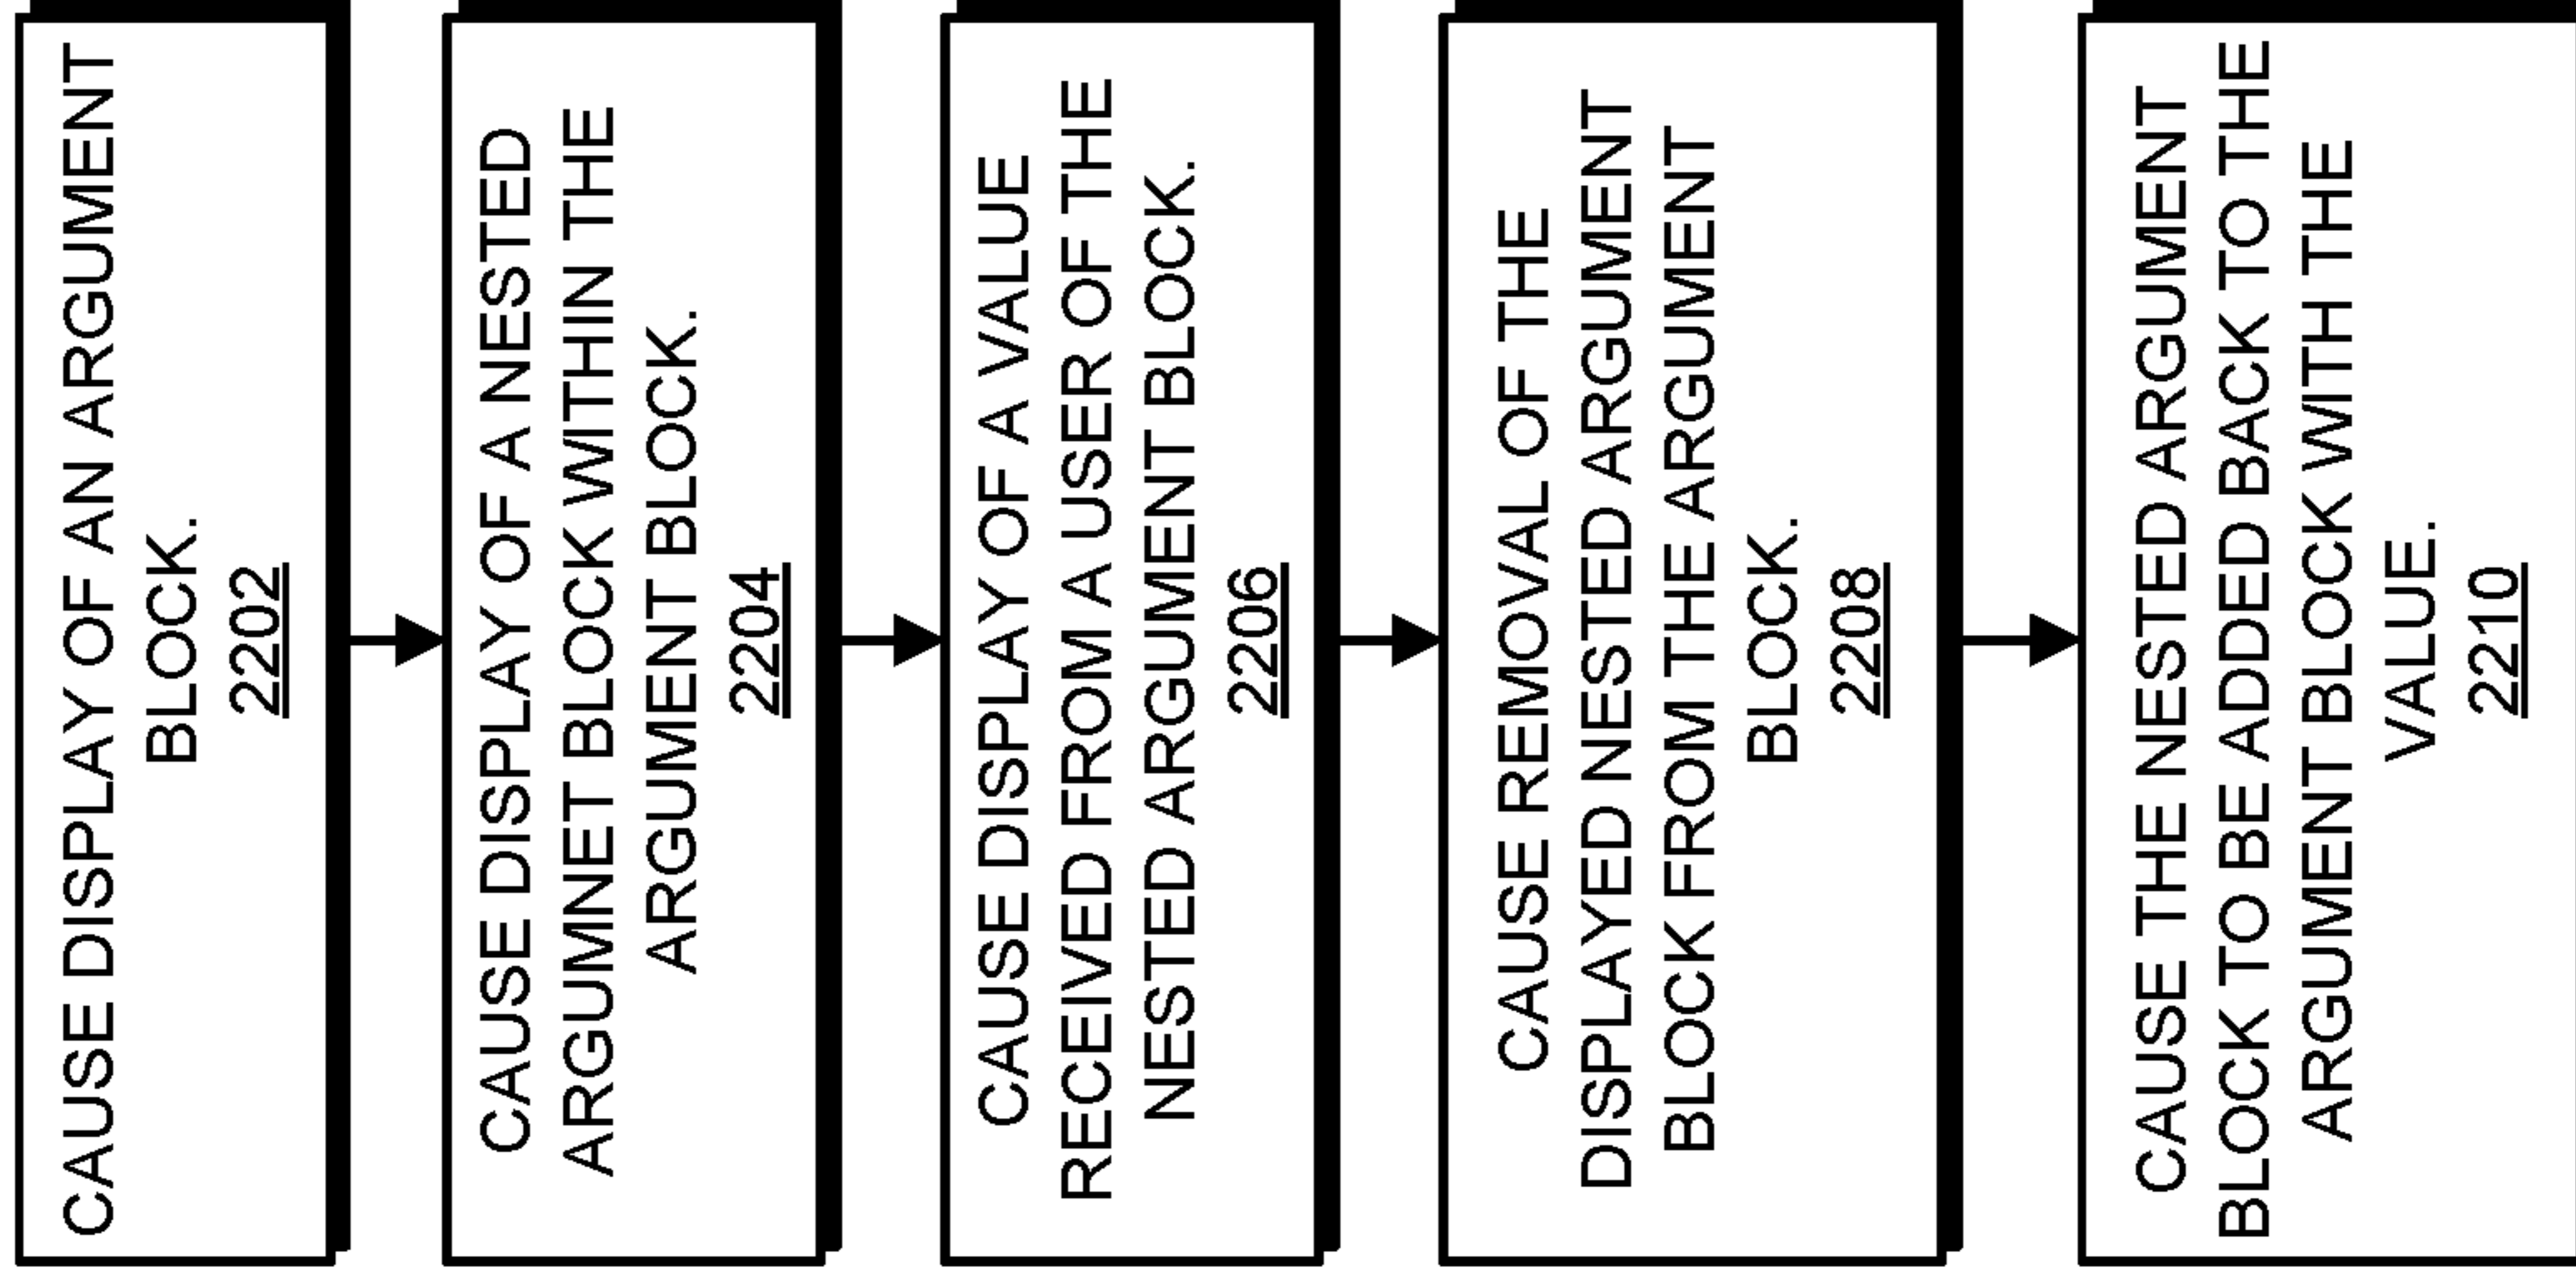


FIG. 20B

**FIG. 21**



**FIG. 22**





## SYNTAX TEMPLATES FOR CODING

## BACKGROUND

Modern data centers often include thousands of hosts that operate collectively to service requests from even larger numbers of remote clients. During operation, components of these data centers can produce significant volumes of machine-generated data. In order to reduce the size of the data, it is typically pre-processed before it is stored. In some instances, the pre-processing includes extracting and storing some of the data, but discarding the remainder of the data. Although this may save storage space in the short term, it can be undesirable in the long term. For example, if the discarded data is later determined to be of use, it may no longer be available.

In some instances, techniques have been developed to apply minimal processing to the data in an attempt to preserve more of the data for later use. For example, the data may be maintained in a relatively unstructured form to reduce the loss of relevant data. Unfortunately, the unstructured nature of much of this data has made it challenging to perform indexing and searching operations because of the difficulty of applying semantic meaning to unstructured data. As the number of hosts and clients associated with a data center continues to grow, processing large volumes of machine-generated data in an intelligent manner and effectively presenting the results of such processing continues to be a priority. Moreover, processing of the data may return a large amount of information that can be difficult for a user to interpret. For example, if a user submits a search of the data, the user may be provided with a large set of search results for the data but may not know how the search results relate to the data itself or how the search results relate to one another. As a result, a user may have a difficult time deciphering what portions of the data or the search results are relevant to her/his inquiry.

## SUMMARY

Embodiments of the present invention are directed to syntax templates for coding.

This summary is provided to introduce a selection of concepts in a simplified form that are further described below in the detailed description. This summary is not intended to identify key features or essential features of the claimed subject matter, nor is it intended to be used in isolation as an aid in determining the scope of the claimed subject matter.

## BRIEF DESCRIPTION OF THE DRAWINGS

Implementations of the present disclosure are described in detail below with reference to the attached drawing figures, wherein:

FIG. 1 illustrates a networked computer environment in which an embodiment may be implemented;

FIG. 2 illustrates a block diagram of an example data intake and query system in which an embodiment may be implemented;

FIG. 3 is a flow diagram that illustrates how indexers process, index, and store data received from forwarders in accordance with the disclosed embodiments;

FIG. 4 is a flow diagram that illustrates how a search head and indexers perform a search query in accordance with the disclosed embodiments;

FIG. 5 illustrates a scenario where a common customer ID is found among log data received from three disparate sources in accordance with the disclosed embodiments;

FIG. 6A illustrates a search screen in accordance with the disclosed embodiments;

FIG. 6B illustrates a data summary dialog that enables a user to select various data sources in accordance with the disclosed embodiments;

FIGS. 7A-7D illustrate a series of user interface screens for an example data model-driven report generation interface in accordance with the disclosed embodiments;

FIG. 8 illustrates an example search query received from a client and executed by search peers in accordance with the disclosed embodiments;

FIG. 9A illustrates a key indicators view in accordance with the disclosed embodiments;

FIG. 9B illustrates an incident review dashboard in accordance with the disclosed embodiments;

FIG. 9C illustrates a proactive monitoring tree in accordance with the disclosed embodiments;

FIG. 9D illustrates a user interface screen displaying both log data and performance data in accordance with the disclosed embodiments;

FIG. 10 illustrates a block diagram of an example cloud-based data intake and query system in which an embodiment may be implemented;

FIG. 11 illustrates a block diagram of an example data intake and query system that performs searches across external data systems in accordance with the disclosed embodiments;

FIGS. 12-14 illustrate a series of user interface screens for an example data model-driven report generation interface in accordance with the disclosed embodiments;

FIGS. 15-17 illustrate example visualizations generated by a reporting application in accordance with the disclosed embodiments;

FIG. 18A illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18B illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18C illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18D illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18E illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18F illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18G illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18H illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18I illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18J illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18K illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18L illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18M illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 18N illustrates a coding screen in accordance with the disclosed embodiments.

FIG. 19A illustrates a syntax definition of a command in accordance with the disclosed embodiments.

FIG. 19B illustrates rules of a command generated from a syntax definition of the command in accordance with the disclosed embodiments.

FIG. 20A illustrates a coded query in accordance with the disclosed embodiments.

FIG. 20B illustrates a reformatted coded query in accordance with the disclosed embodiments.

FIG. 21 presents a flowchart illustrating a method in accordance with the disclosed embodiments.

FIG. 22 presents a flowchart illustrating a method in accordance with the disclosed embodiments.

## DETAILED DESCRIPTION

Embodiments are described herein according to the following outline:

- 1.0. General Overview
- 2.0. Operating Environment
  - 2.1. Host Devices
  - 2.2. Client Devices
  - 2.3. Client Device Applications
  - 2.4. Data Server System
  - 2.5. Data Ingestion
    - 2.5.1. Input
    - 2.5.2. Parsing
    - 2.5.3. Indexing
  - 2.6. Query Processing
  - 2.7. Field Extraction
  - 2.8. Example Search Screen
  - 2.9. Data Modelling
  - 2.10. Acceleration Techniques
    - 2.10.1. Aggregation Technique
    - 2.10.2. Keyword Index
    - 2.10.3. High Performance Analytics Store
    - 2.10.4. Accelerating Report Generation
  - 2.11. Security Features
  - 2.12. Data Center Monitoring
  - 2.13. Cloud-Based System Overview
  - 2.14. Searching Externally Archived Data
    - 2.14.1. ERP Process Features
  - 2.15. IT Service Monitoring
- 3.0. Code Completion
  - 3.1. Example Coding Interface
  - 3.2. Contextual Suggestions and Syntax Templates
  - 3.3. Rule Generation
  - 3.4. Code Reformatting
  - 3.5. Additional Implementations

### 1.0. GENERAL OVERVIEW

Modern data centers and other computing environments can comprise anywhere from a few host computer systems to thousands of systems configured to process data, service requests from remote clients, and perform numerous other computational tasks. During operation, various components within these computing environments often generate significant volumes of machine-generated data. For example, machine data is generated by various components in the information technology (IT) environments, such as servers, sensors, routers, mobile devices, Internet of Things (IoT) devices, etc. Machine-generated data can include system logs, network packet data, sensor data, application program data, error logs, stack traces, system performance data, etc. In general, machine-generated data can also include performance data, diagnostic information, and many other types of data that can be analyzed to diagnose performance problems, monitor user interactions, and to derive other insights.

A number of tools are available to analyze machine data, that is, machine-generated data. In order to reduce the size of the potentially vast amount of machine data that may be generated, many of these tools typically pre-process the data based on anticipated data-analysis needs. For example, pre-specified data items may be extracted from the machine data and stored in a database to facilitate efficient retrieval and analysis of those data items at search time. However, the rest of the machine data typically is not saved and discarded during pre-processing. As storage capacity becomes progressively cheaper and more plentiful, there are fewer incentives to discard these portions of machine data and many reasons to retain more of the data.

This plentiful storage capacity is presently making it feasible to store massive quantities of minimally processed machine data for later retrieval and analysis. In general, storing minimally processed machine data and performing analysis operations at search time can provide greater flexibility because it enables an analyst to search all of the machine data, instead of searching only a pre-specified set of data items. This may enable an analyst to investigate different aspects of the machine data that previously were unavailable for analysis.

However, analyzing and searching massive quantities of machine data presents a number of challenges. For example, a data center, servers, or network appliances may generate many different types and formats of machine data (e.g., system logs, network packet data (e.g., wire data, etc.), sensor data, application program data, error logs, stack traces, system performance data, operating system data, virtualization data, etc.) from thousands of different components, which can collectively be very time-consuming to analyze. In another example, mobile devices may generate large amounts of information relating to data accesses, application performance, operating system performance, network performance, etc. There can be millions of mobile devices that report these types of information.

These challenges can be addressed by using an event-based data intake and query system, such as the SPLUNK® ENTERPRISE system developed by Splunk Inc. of San Francisco, Calif. The SPLUNK® ENTERPRISE system is the leading platform for providing real-time operational intelligence that enables organizations to collect, index, and search machine-generated data from various websites, applications, servers, networks, and mobile devices that power their businesses. The SPLUNK® ENTERPRISE system is particularly useful for analyzing data which is commonly found in system log files, network data, and other data input sources. Although many of the techniques described herein are explained with reference to a data intake and query system similar to the SPLUNK® ENTERPRISE system, these techniques are also applicable to other types of data systems.

In the SPLUNK® ENTERPRISE system, machine-generated data are collected and stored as “events”. An event comprises a portion of the machine-generated data and is associated with a specific point in time. For example, events may be derived from “time series data,” where the time series data comprises a sequence of data points (e.g., performance measurements from a computer system, etc.) that are associated with successive points in time. In general, each event can be associated with a timestamp that is derived from the raw data in the event, determined through interpolation between temporally proximate events having known timestamps, or determined based on other configurable rules for associating timestamps with events, etc.

In some instances, machine data can have a predefined format, where data items with specific data formats are stored at predefined locations in the data. For example, the machine data may include data stored as fields in a database table. In other instances, machine data may not have a predefined format, that is, the data is not at fixed, predefined locations, but the data does have repeatable patterns and is not random. This means that some machine data can comprise various data items of different data types and that may be stored at different locations within the data. For example, when the data source is an operating system log, an event can include one or more lines from the operating system log containing raw data that includes different types of performance and diagnostic information associated with a specific point in time.

Examples of components which may generate machine data from which events can be derived include, but are not limited to, web servers, application servers, databases, firewalls, routers, operating systems, and software applications that execute on computer systems, mobile devices, sensors, Internet of Things (IoT) devices, etc. The data generated by such data sources can include, for example and without limitation, server log files, activity log files, configuration files, messages, network packet data, performance measurements, sensor measurements, etc.

The SPLUNK® ENTERPRISE system uses flexible schema to specify how to extract information from the event data. A flexible schema may be developed and redefined as needed. Note that a flexible schema may be applied to event data “on the fly,” when it is needed (e.g., at search time, index time, ingestion time, etc.). When the schema is not applied to event data until search time it may be referred to as a “late-binding schema.”

During operation, the SPLUNK® ENTERPRISE system starts with raw input data (e.g., one or more system logs, streams of network packet data, sensor data, application program data, error logs, stack traces, system performance data, etc.). The system divides this raw data into blocks (e.g., buckets of data, each associated with a specific time frame, etc.), and parses the raw data to produce timestamped events. The system stores the timestamped events in a data store. The system enables users to run queries against the stored data to, for example, retrieve events that meet criteria specified in a query, such as containing certain keywords or having specific values in defined fields. As used herein throughout, data that is part of an event is referred to as “event data”. In this context, the term “field” refers to a location in the event data containing one or more values for a specific data item. As will be described in more detail herein, the fields are defined by extraction rules (e.g., regular expressions) that derive one or more values from the portion of raw machine data in each event that has a particular field specified by an extraction rule. The set of values so produced are semantically-related (such as IP address), even though the raw machine data in each event may be in different formats (e.g., semantically-related values may be in different positions in the events derived from different sources).

As noted above, the SPLUNK® ENTERPRISE system utilizes a late-binding schema to event data while performing queries on events. One aspect of a late-binding schema is applying “extraction rules” to event data to extract values for specific fields during search time. More specifically, the extraction rules for a field can include one or more instructions that specify how to extract a value for the field from the event data. An extraction rule can generally include any type of instruction for extracting values from data in events. In some cases, an extraction rule comprises a regular expres-

sion where a sequence of characters form a search pattern, in which case the rule is referred to as a “regex rule.” The system applies the regex rule to the event data to extract values for associated fields in the event data by searching the event data for the sequence of characters defined in the regex rule.

In the SPLUNK® ENTERPRISE system, a field extractor may be configured to automatically generate extraction rules for certain field values in the events when the events are being created, indexed, or stored, or possibly at a later time. Alternatively, a user may manually define extraction rules for fields using a variety of techniques. In contrast to a conventional schema for a database system, a late-binding schema is not defined at data ingestion time. Instead, the late-binding schema can be developed on an ongoing basis until the time a query is actually executed. This means that extraction rules for the fields in a query may be provided in the query itself, or may be located during execution of the query. Hence, as a user learns more about the data in the events, the user can continue to refine the late-binding schema by adding new fields, deleting fields, or modifying the field extraction rules for use the next time the schema is used by the system. Because the SPLUNK® ENTERPRISE system maintains the underlying raw data and uses late-binding schema for searching the raw data, it enables a user to continue investigating and learn valuable insights about the raw data.

In some embodiments, a common field name may be used to reference two or more fields containing equivalent data items, even though the fields may be associated with different types of events that possibly have different data formats and different extraction rules. By enabling a common field name to be used to identify equivalent fields from different types of events generated by disparate data sources, the system facilitates use of a “common information model” (CIM) across the disparate data sources (further discussed with respect to FIG. 5).

## 2.0. OPERATING ENVIRONMENT

FIG. 1 illustrates a networked computer system **100** in which an embodiment may be implemented. Those skilled in the art would understand that FIG. 1 represents one example of a networked computer system and other embodiments may use different arrangements.

The networked computer system **100** comprises one or more computing devices. These one or more computing devices comprise any combination of hardware and software configured to implement the various logical components described herein. For example, the one or more computing devices may include one or more memories that store instructions for implementing the various components described herein, one or more hardware processors configured to execute the instructions stored in the one or more memories, and various data repositories in the one or more memories for storing data structures utilized and manipulated by the various components.

In an embodiment, one or more client devices **102** are coupled to one or more host devices **106** and a data intake and query system **108** via one or more networks **104**. Networks **104** broadly represent one or more LANs, WANs, cellular networks (e.g., LTE, HSPA, 3G and other cellular technologies), and/or networks using any of wired, wireless, terrestrial microwave, or satellite links, and may include the public Internet.

### 2.1. Host Devices

In the illustrated embodiment, a system **100** includes one or more host devices **106**. Host devices **106** may broadly

include any number of computers, virtual machine instances, and/or data centers that are configured to host or execute one or more instances of host applications **114**. In general, a host device **106** may be involved, directly or indirectly, in processing requests received from client devices **102**. Each host device **106** may comprise, for example, one or more of a network device, a web server, an application server, a database server, etc. A collection of host devices **106** may be configured to implement a network-based service. For example, a provider of a network-based service may configure one or more host devices **106** and host applications **114** (e.g., one or more web servers, application servers, database servers, etc.) to collectively implement the network-based application.

In general, client devices **102** communicate with one or more host applications **114** to exchange information. The communication between a client device **102** and a host application **114** may, for example, be based on the Hypertext Transfer Protocol (HTTP) or any other network protocol. Content delivered from the host application **114** to a client device **102** may include, for example, HTML documents, media content, etc. The communication between a client device **102** and host application **114** may include sending various requests and receiving data packets. For example, in general, a client device **102** or application running on a client device may initiate communication with a host application **114** by making a request for a specific resource (e.g., based on an HTTP request), and the application server may respond with the requested content stored in one or more response packets.

In the illustrated embodiment, one or more of host applications **114** may generate various types of performance data during operation, including event logs, network data, sensor data, and other types of machine-generated data. For example, a host application **114** comprising a web server may generate one or more web server logs in which details of interactions between the web server and any number of client devices **102** is recorded. As another example, a host device **106** comprising a router may generate one or more router logs that record information related to network traffic managed by the router. As yet another example, a host application **114** comprising a database server may generate one or more logs that record information related to requests sent from other host applications **114** (e.g., web servers or application servers) for data managed by the database server.

## 2.2. Client Devices

Client devices **102** of FIG. 1 represent any computing device capable of interacting with one or more host devices **106** via a network **104**. Examples of client devices **102** may include, without limitation, smart phones, tablet computers, handheld computers, wearable devices, laptop computers, desktop computers, servers, portable media players, gaming devices, and so forth. In general, a client device **102** can provide access to different content, for instance, content provided by one or more host devices **106**, etc. Each client device **102** may comprise one or more client applications **110**, described in more detail in a separate section hereinafter.

## 2.3. Client Device Applications

In an embodiment, each client device **102** may host or execute one or more client applications **110** that are capable of interacting with one or more host devices **106** via one or more networks **104**. For instance, a client application **110**

may be or comprise a web browser that a user may use to navigate to one or more websites or other resources provided by one or more host devices **106**. As another example, a client application **110** may comprise a mobile application or “app.” For example, an operator of a network-based service hosted by one or more host devices **106** may make available one or more mobile apps that enable users of client devices **102** to access various resources of the network-based service. As yet another example, client applications **110** may include background processes that perform various operations without direct interaction from a user. A client application **110** may include a “plug-in” or “extension” to another application, such as a web browser plug-in or extension.

In an embodiment, a client application **110** may include a monitoring component **112**. At a high level, the monitoring component **112** comprises a software component or other logic that facilitates generating performance data related to a client device’s operating state, including monitoring network traffic sent and received from the client device and collecting other device and/or application-specific information. Monitoring component **112** may be an integrated component of a client application **110**, a plug-in, an extension, or any other type of add-on component. Monitoring component **112** may also be a stand-alone process.

In one embodiment, a monitoring component **112** may be created when a client application **110** is developed, for example, by an application developer using a software development kit (SDK). The SDK may include custom monitoring code that can be incorporated into the code implementing a client application **110**. When the code is converted to an executable application, the custom code implementing the monitoring functionality can become part of the application itself.

In some cases, an SDK or other code for implementing the monitoring functionality may be offered by a provider of a data intake and query system, such as a system **108**. In such cases, the provider of the system **108** can implement the custom code so that performance data generated by the monitoring functionality is sent to the system **108** to facilitate analysis of the performance data by a developer of the client application or other users.

In an embodiment, the custom monitoring code may be incorporated into the code of a client application **110** in a number of different ways, such as the insertion of one or more lines in the client application code that call or otherwise invoke the monitoring component **112**. As such, a developer of a client application **110** can add one or more lines of code into the client application **110** to trigger the monitoring component **112** at desired points during execution of the application. Code that triggers the monitoring component may be referred to as a monitor trigger. For instance, a monitor trigger may be included at or near the beginning of the executable code of the client application **110** such that the monitoring component **112** is initiated or triggered as the application is launched, or included at other points in the code that correspond to various actions of the client application, such as sending a network request or displaying a particular interface.

In an embodiment, the monitoring component **112** may monitor one or more aspects of network traffic sent and/or received by a client application **110**. For example, the monitoring component **112** may be configured to monitor data packets transmitted to and/or from one or more host applications **114**. Incoming and/or outgoing data packets can be read or examined to identify network data contained within the packets, for example, and other aspects of data packets can be analyzed to determine a number of network

performance statistics. Monitoring network traffic may enable information to be gathered particular to the network performance associated with a client application **110** or set of applications.

In an embodiment, network performance data refers to any type of data that indicates information about the network and/or network performance. Network performance data may include, for instance, a URL requested, a connection type (e.g., HTTP, HTTPS, etc.), a connection start time, a connection end time, an HTTP status code, request length, response length, request headers, response headers, connection status (e.g., completion, response time(s), failure, etc.), and the like. Upon obtaining network performance data indicating performance of the network, the network performance data can be transmitted to a data intake and query system **108** for analysis.

Upon developing a client application **110** that incorporates a monitoring component **112**, the client application **110** can be distributed to client devices **102**. Applications generally can be distributed to client devices **102** in any manner, or they can be pre-loaded. In some cases, the application may be distributed to a client device **102** via an application marketplace or other application distribution system. For instance, an application marketplace or other application distribution system might distribute the application to a client device based on a request from the client device to download the application.

Examples of functionality that enables monitoring performance of a client device are described in U.S. patent application Ser. No. 14/524,748, entitled “UTILIZING PACKET HEADERS TO MONITOR NETWORK TRAFFIC IN ASSOCIATION WITH A CLIENT DEVICE”, filed on 27 Oct. 2014, and which is hereby incorporated by reference in its entirety for all purposes.

In an embodiment, the monitoring component **112** may also monitor and collect performance data related to one or more aspects of the operational state of a client application **110** and/or client device **102**. For example, a monitoring component **112** may be configured to collect device performance information by monitoring one or more client device operations, or by making calls to an operating system and/or one or more other applications executing on a client device **102** for performance information. Device performance information may include, for instance, a current wireless signal strength of the device, a current connection type and network carrier, current memory performance information, a geographic location of the device, a device orientation, and any other information related to the operational state of the client device.

In an embodiment, the monitoring component **112** may also monitor and collect other device profile information including, for example, a type of client device, a manufacturer and model of the device, versions of various software applications installed on the device, and so forth.

In general, a monitoring component **112** may be configured to generate performance data in response to a monitor trigger in the code of a client application **110** or other triggering application event, as described above, and to store the performance data in one or more data records. Each data record, for example, may include a collection of field-value pairs, each field-value pair storing a particular item of performance data in association with a field for the item. For example, a data record generated by a monitoring component **112** may include a “networkLatency” field (not shown in the Figure) in which a value is stored. This field indicates a network latency measurement associated with one or more network requests. The data record may include a “state” field

to store a value indicating a state of a network connection, and so forth for any number of aspects of collected performance data.

#### 2.4. Data Server System

FIG. 2 depicts a block diagram of an exemplary data intake and query system **108**, similar to the SPLUNK® ENTERPRISE system. System **108** includes one or more forwarders **204** that receive data from a variety of input data sources **202**, and one or more indexers **206** that process and store the data in one or more data stores **208**. These forwarders and indexers can comprise separate computer systems, or may alternatively comprise separate processes executing on one or more computer systems.

Each data source **202** broadly represents a distinct source of data that can be consumed by a system **108**. Examples of a data source **202** include, without limitation, data files, directories of files, data sent over a network, event logs, registries, etc.

During operation, the forwarders **204** identify which indexers **206** receive data collected from a data source **202** and forward the data to the appropriate indexers. Forwarders **204** can also perform operations on the data before forwarding, including removing extraneous data, detecting timestamps in the data, parsing data, indexing data, routing data based on criteria relating to the data being routed, and/or performing other data transformations.

In an embodiment, a forwarder **204** may comprise a service accessible to client devices **102** and host devices **106** via a network **104**. For example, one type of forwarder **204** may be capable of consuming vast amounts of real-time data from a potentially large number of client devices **102** and/or host devices **106**. The forwarder **204** may, for example, comprise a computing device which implements multiple data pipelines or “queues” to handle forwarding of network data to indexers **206**. A forwarder **204** may also perform many of the functions that are performed by an indexer. For example, a forwarder **204** may perform keyword extractions on raw data or parse raw data to create events. A forwarder **204** may generate time stamps for events. Additionally or alternatively, a forwarder **204** may perform routing of events to indexers. Data store **208** may contain events derived from machine data from a variety of sources all pertaining to the same component in an IT environment, and this data may be produced by the machine in question or by other components in the IT environment.

#### 2.5. Data Ingestion

FIG. 3 depicts a flow chart illustrating an example data flow performed by Data Intake and Query system **108**, in accordance with the disclosed embodiments. The data flow illustrated in FIG. 3 is provided for illustrative purposes only; those skilled in the art would understand that one or more of the steps of the processes illustrated in FIG. 3 may be removed or the ordering of the steps may be changed. Furthermore, for the purposes of illustrating a clear example, one or more particular system components are described in the context of performing various operations during each of the data flow stages. For example, a forwarder is described as receiving and processing data during an input phase; an indexer is described as parsing and indexing data during parsing and indexing phases; and a search head is described as performing a search query during a search phase. How-

ever, other system arrangements and distributions of the processing steps across system components may be used.

#### 2.5.1. Input

At block **302**, a forwarder receives data from an input source, such as a data source **202** shown in FIG. **2**. A forwarder initially may receive the data as a raw data stream generated by the input source. For example, a forwarder may receive a data stream from a log file generated by an application server, from a stream of network data from a network device, or from any other source of data. In one embodiment, a forwarder receives the raw data and may segment the data stream into “blocks”, or “buckets,” possibly of a uniform data size, to facilitate subsequent processing steps.

At block **304**, a forwarder or other system component annotates each block generated from the raw data with one or more metadata fields. These metadata fields may, for example, provide information related to the data block as a whole and may apply to each event that is subsequently derived from the data in the data block. For example, the metadata fields may include separate fields specifying each of a host, a source, and a source type related to the data block. A host field may contain a value identifying a host name or IP address of a device that generated the data. A source field may contain a value identifying a source of the data, such as a pathname of a file or a protocol and port related to received network data. A source type field may contain a value specifying a particular source type label for the data. Additional metadata fields may also be included during the input phase, such as a character encoding of the data, if known, and possibly other values that provide information relevant to later processing steps. In an embodiment, a forwarder forwards the annotated data blocks to another system component (typically an indexer) for further processing.

The SPLUNK® ENTERPRISE system allows forwarding of data from one SPLUNK® ENTERPRISE instance to another, or even to a third-party system. SPLUNK® ENTERPRISE system can employ different types of forwarders in a configuration.

In an embodiment, a forwarder may contain the essential components needed to forward data. It can gather data from a variety of inputs and forward the data to a SPLUNK® ENTERPRISE server for indexing and searching. It also can tag metadata (e.g., source, source type, host, etc.).

Additionally or optionally, in an embodiment, a forwarder has the capabilities of the aforementioned forwarder as well as additional capabilities. The forwarder can parse data before forwarding the data (e.g., associate a time stamp with a portion of data and create an event, etc.) and can route data based on criteria such as source or type of event. It can also index data locally while forwarding the data to another indexer.

#### 2.5.2. Parsing

At block **306**, an indexer receives data blocks from a forwarder and parses the data to organize the data into events. In an embodiment, to organize the data into events, an indexer may determine a source type associated with each data block (e.g., by extracting a source type label from the metadata fields associated with the data block, etc.) and refer to a source type configuration corresponding to the identified source type. The source type definition may include one or more properties that indicate to the indexer to automatically

determine the boundaries of events within the data. In general, these properties may include regular expression-based rules or delimiter rules where, for example, event boundaries may be indicated by predefined characters or character strings. These predefined characters may include punctuation marks or other special characters including, for example, carriage returns, tabs, spaces, line breaks, etc. If a source type for the data is unknown to the indexer, an indexer may infer a source type for the data by examining the structure of the data. Then, it can apply an inferred source type definition to the data to create the events.

At block **308**, the indexer determines a timestamp for each event. Similar to the process for creating events, an indexer may again refer to a source type definition associated with the data to locate one or more properties that indicate instructions for determining a timestamp for each event. The properties may, for example, instruct an indexer to extract a time value from a portion of data in the event, to interpolate time values based on timestamps associated with temporally proximate events, to create a timestamp based on a time the event data was received or generated, to use the timestamp of a previous event, or use any other rules for determining timestamps.

At block **310**, the indexer associates with each event one or more metadata fields including a field containing the timestamp (in some embodiments, a timestamp may be included in the metadata fields) determined for the event. These metadata fields may include a number of “default fields” that are associated with all events, and may also include one more custom fields as defined by a user. Similar to the metadata fields associated with the data blocks at block **304**, the default metadata fields associated with each event may include a host, source, and source type field including or in addition to a field storing the timestamp.

At block **312**, an indexer may optionally apply one or more transformations to data included in the events created at block **306**. For example, such transformations can include removing a portion of an event (e.g., a portion used to define event boundaries, extraneous characters from the event, other extraneous text, etc.), masking a portion of an event (e.g., masking a credit card number), removing redundant portions of an event, etc. The transformations applied to event data may, for example, be specified in one or more configuration files and referenced by one or more source type definitions.

#### 2.5.3. Indexing

At blocks **314** and **316**, an indexer can optionally generate a keyword index to facilitate fast keyword searching for event data. To build a keyword index, at block **314**, the indexer identifies a set of keywords in each event. At block **316**, the indexer includes the identified keywords in an index, which associates each stored keyword with reference pointers to events containing that keyword (or to locations within events where that keyword is located, other location identifiers, etc.). When an indexer subsequently receives a keyword-based query, the indexer can access the keyword index to quickly identify events containing the keyword.

In some embodiments, the keyword index may include entries for name-value pairs found in events, where a name-value pair can include a pair of keywords connected by a symbol, such as an equals sign or colon. This way, events containing these name-value pairs can be quickly located. In some embodiments, fields can automatically be generated for some or all of the name-value pairs at the time of indexing. For example, if the string “dest=10.0.1.2” is

found in an event, a field named “dest” may be created for the event, and assigned a value of “10.0.1.2”.

At block **318**, the indexer stores the events with an associated timestamp in a data store **208**. Timestamps enable a user to search for events based on a time range. In one embodiment, the stored events are organized into “buckets,” where each bucket stores events associated with a specific time range based on the timestamps associated with each event. This may not only improve time-based searching, but also allows for events with recent timestamps, which may have a higher likelihood of being accessed, to be stored in a faster memory to facilitate faster retrieval. For example, buckets containing the most recent events can be stored in flash memory rather than on a hard disk.

Each indexer **206** may be responsible for storing and searching a subset of the events contained in a corresponding data store **208**. By distributing events among the indexers and data stores, the indexers can analyze events for a query in parallel. For example, using map-reduce techniques, each indexer returns partial responses for a subset of events to a search head that combines the results to produce an answer for the query. By storing events in buckets for specific time ranges, an indexer may further optimize data retrieval process by searching buckets corresponding to time ranges that are relevant to a query.

Moreover, events and buckets can also be replicated across different indexers and data stores to facilitate high availability and disaster recovery as described in U.S. patent application Ser. No. 14/266,812, entitled “SITE-BASED SEARCH AFFINITY”, filed on 30 Apr. 2014, and in U.S. patent application Ser. No. 14/266,817, entitled “MULTI-SITE CLUSTERING”, also filed on 30 Apr. 2014, each of which is hereby incorporated by reference in its entirety for all purposes.

## 2.6. Query Processing

FIG. **4** is a flow diagram that illustrates an exemplary process that a search head and one or more indexers may perform during a search query. At block **402**, a search head receives a search query from a client. At block **404**, the search head analyzes the search query to determine what portion(s) of the query can be delegated to indexers and what portions of the query can be executed locally by the search head. At block **406**, the search head distributes the determined portions of the query to the appropriate indexers. In an embodiment, a search head cluster may take the place of an independent search head where each search head in the search head cluster coordinates with peer search heads in the search head cluster to schedule jobs, replicate search results, update configurations, fulfill search requests, etc. In an embodiment, the search head (or each search head) communicates with a master node (also known as a cluster master, not shown in Fig.) that provides the search head with a list of indexers to which the search head can distribute the determined portions of the query. The master node maintains a list of active indexers and can also designate which indexers may have responsibility for responding to queries over certain sets of events. A search head may communicate with the master node before the search head distributes queries to indexers to discover the addresses of active indexers.

At block **408**, the indexers to which the query was distributed, search data stores associated with them for events that are responsive to the query. To determine which events are responsive to the query, the indexer searches for events that match the criteria specified in the query. These

criteria can include matching keywords or specific values for certain fields. The searching operations at block **408** may use the late-binding schema to extract values for specified fields from events at the time the query is processed. In an embodiment, one or more rules for extracting field values may be specified as part of a source type definition. The indexers may then either send the relevant events back to the search head, or use the events to determine a partial result, and send the partial result back to the search head.

At block **410**, the search head combines the partial results and/or events received from the indexers to produce a final result for the query. This final result may comprise different types of data depending on what the query requested. For example, the results can include a listing of matching events returned by the query, or some type of visualization of the data from the returned events. In another example, the final result can include one or more calculated values derived from the matching events.

The results generated by the system **108** can be returned to a client using different techniques. For example, one technique streams results or relevant events back to a client in real-time as they are identified. Another technique waits to report the results to the client until a complete set of results (which may include a set of relevant events or a result based on relevant events) is ready to return to the client. Yet another technique streams interim results or relevant events back to the client in real-time until a complete set of results is ready, and then returns the complete set of results to the client. In another technique, certain results are stored as “search jobs” and the client may retrieve the results by referring the search jobs.

The search head can also perform various operations to make the search more efficient. For example, before the search head begins execution of a query, the search head can determine a time range for the query and a set of common keywords that all matching events include. The search head may then use these parameters to query the indexers to obtain a superset of the eventual results. Then, during a filtering stage, the search head can perform field-extraction operations on the superset to produce a reduced set of search results. This speeds up queries that are performed on a periodic basis.

## 2.7. Field Extraction

The search head **210** allows users to search and visualize event data extracted from raw machine data received from homogenous data sources. It also allows users to search and visualize event data extracted from raw machine data received from heterogeneous data sources. The search head **210** includes various mechanisms, which may additionally reside in an indexer **206**, for processing a query. Splunk Processing Language (SPL), used in conjunction with the SPLUNK® ENTERPRISE system, can be utilized to make a query. SPL is a pipelined search language in which a set of inputs is operated on by a first command in a command line, and then a subsequent command following the pipe symbol “|” operates on the results produced by the first command, and so on for additional commands. Other query languages, such as the Structured Query Language (“SQL”), can be used to create a query.

In response to receiving the search query, search head **210** uses extraction rules to extract values for the fields associated with a field or fields in the event data being searched. The search head **210** obtains extraction rules that specify how to extract a value for certain fields from an event. Extraction rules can comprise regex rules that specify how

to extract values for the relevant fields. In addition to specifying how to extract field values, the extraction rules may also include instructions for deriving a field value by performing a function on a character string or value retrieved by the extraction rule. For example, a transformation rule may truncate a character string, or convert the character string into a different data format. In some cases, the query itself can specify one or more extraction rules.

The search head **210** can apply the extraction rules to event data that it receives from indexers **206**. Indexers **206** may apply the extraction rules to events in an associated data store **208**. Extraction rules can be applied to all the events in a data store, or to a subset of the events that have been filtered based on some criteria (e.g., event time stamp values, etc.). Extraction rules can be used to extract one or more values for a field from events by parsing the event data and examining the event data for one or more patterns of characters, numbers, delimiters, etc., that indicate where the field begins and, optionally, ends.

FIG. **5** illustrates an example of raw machine data received from disparate data sources. In this example, a user submits an order for merchandise using a vendor's shopping application program **501** running on the user's system. In this example, the order was not delivered to the vendor's server due to a resource exception at the destination server that is detected by the middleware code **502**. The user then sends a message to the customer support **503** to complain about the order failing to complete. The three systems **501**, **502**, and **503** are disparate systems that do not have a common logging format. The order application **501** sends log data **504** to the SPLUNK® ENTERPRISE system in one format, the middleware code **502** sends error log data **505** in a second format, and the support server **503** sends log data **506** in a third format.

Using the log data received at one or more indexers **206** from the three systems the vendor can uniquely obtain an insight into user activity, user experience, and system behavior. The search head **210** allows the vendor's administrator to search the log data from the three systems that one or more indexers **206** are responsible for searching, thereby obtaining correlated information, such as the order number and corresponding customer ID number of the person placing the order. The system also allows the administrator to see a visualization of related events via a user interface. The administrator can query the search head **210** for customer ID field value matches across the log data from the three systems that are stored at the one or more indexers **206**. The customer ID field value exists in the data gathered from the three systems, but the customer ID field value may be located in different areas of the data given differences in the architecture of the systems—there is a semantic relationship between the customer ID field values generated by the three systems. The search head **210** requests event data from the one or more indexers **206** to gather relevant event data from the three systems. It then applies extraction rules to the event data in order to extract field values that it can correlate. The search head may apply a different extraction rule to each set of events from each system when the event data format differs among systems. In this example, the user interface can display to the administrator the event data corresponding to the common customer ID field values **507**, **508**, and **509**, thereby providing the administrator with insight into a customer's experience.

Note that query results can be returned to a client, a search head, or any other system component for further processing. In general, query results may include a set of one or more events, a set of one or more values obtained from the events,

a subset of the values, statistics calculated based on the values, a report containing the values, or a visualization, such as a graph or chart, generated from the values.

## 2.8. Example Search Screen

FIG. **6A** illustrates an example search screen **600** in accordance with the disclosed embodiments. Search screen **600** includes a search bar **602** that accepts user input in the form of a search string. It also includes a time range picker **612** that enables the user to specify a time range for the search. For “historical searches” the user can select a specific time range, or alternatively a relative time range, such as “today,” “yesterday” or “last week.” For “real-time searches,” the user can select the size of a preceding time window to search for real-time events. Search screen **600** also initially displays a “data summary” dialog as is illustrated in FIG. **6B** that enables the user to select different sources for the event data, such as by selecting specific hosts and log files.

After the search is executed, the search screen **600** in FIG. **6A** can display the results through search results tabs **604**, wherein search results tabs **604** includes: an “events tab” that displays various information about events returned by the search; a “statistics tab” that displays statistics about the search results; and a “visualization tab” that displays various visualizations of the search results. The events tab illustrated in FIG. **6A** displays a timeline graph **605** that graphically illustrates the number of events that occurred in one-hour intervals over the selected time range. It also displays an events list **608** that enables a user to view the raw data in each of the returned events. It additionally displays a fields sidebar **606** that includes statistics about occurrences of specific fields in the returned events, including “selected fields” that are pre-selected by the user, and “interesting fields” that are automatically selected by the system based on pre-specified criteria.

## 2.9. Data Models

A data model is a hierarchically structured search-time mapping of semantic knowledge about one or more datasets. It encodes the domain knowledge necessary to build a variety of specialized searches of those datasets. Those searches, in turn, can be used to generate reports.

A data model is composed of one or more “objects” (or “data model objects”) that define or otherwise correspond to a specific set of data.

Objects in data models can be arranged hierarchically in parent/child relationships. Each child object represents a subset of the dataset covered by its parent object. The top-level objects in data models are collectively referred to as “root objects.”

Child objects have inheritance. Data model objects are defined by characteristics that mostly break down into constraints and attributes. Child objects inherit constraints and attributes from their parent objects and have additional constraints and attributes of their own. Child objects provide a way of filtering events from parent objects. Because a child object always provides an additional constraint in addition to the constraints it has inherited from its parent object, the dataset it represents is always a subset of the dataset that its parent represents.

For example, a first data model object may define a broad set of data pertaining to e-mail activity generally, and another data model object may define specific datasets within the broad dataset, such as a subset of the e-mail data



pertaining specifically to e-mails sent. Examples of data models can include electronic mail, authentication, databases, intrusion detection, malware, application state, alerts, compute inventory, network sessions, network traffic, performance, audits, updates, vulnerabilities, etc. Data models and their objects can be designed by knowledge managers in an organization, and they can enable downstream users to quickly focus on a specific set of data. For example, a user can simply select an “e-mail activity” data model object to access a dataset relating to e-mails generally (e.g., sent or received), or select an “e-mails sent” data model object (or data sub-model object) to access a dataset relating to e-mails sent.

A data model object may be defined by (1) a set of search constraints, and (2) a set of fields. Thus, a data model object can be used to quickly search data to identify a set of events and to identify a set of fields to be associated with the set of events. For example, an “e-mails sent” data model object may specify a search for events relating to e-mails that have been sent, and specify a set of fields that are associated with the events. Thus, a user can retrieve and use the “e-mails sent” data model object to quickly search source data for events relating to sent e-mails, and may be provided with a listing of the set of fields relevant to the events in a user interface screen.

A child of the parent data model may be defined by a search (typically a narrower search) that produces a subset of the events that would be produced by the parent data model’s search. The child’s set of fields can include a subset of the set of fields of the parent data model and/or additional fields. Data model objects that reference the subsets can be arranged in a hierarchical manner, so that child subsets of events are proper subsets of their parents. A user iteratively applies a model development tool (not shown in Fig.) to prepare a query that defines a subset of events and assigns an object name to that subset. A child subset is created by further limiting a query that generated a parent subset. A late-binding schema of field extraction rules is associated with each object or subset in the data model.

Data definitions in associated schemas can be taken from the common information model (CIM) or can be devised for a particular schema and optionally added to the CIM. Child objects inherit fields from parents and can include fields not present in parents. A model developer can select fewer extraction rules than are available for the sources returned by the query that defines events belonging to a model. Selecting a limited set of extraction rules can be a tool for simplifying and focusing the data model, while allowing a user flexibility to explore the data subset. Development of a data model is further explained in U.S. Pat. Nos. 8,788,525 and 8,788,526, both entitled “DATA MODEL FOR MACHINE DATA FOR SEMANTIC SEARCH”, both issued on 22 Jul. 2014, U.S. Pat. No. 8,983,994, entitled “GENERATION OF A DATA MODEL FOR SEARCHING MACHINE DATA”, issued on 17 Mar. 2015, U.S. patent application Ser. No. 14/611,232, entitled “GENERATION OF A DATA MODEL APPLIED TO QUERIES”, filed on 31 Jan. 2015, and U.S. patent application Ser. No. 14/815,884, entitled “GENERATION OF A DATA MODEL APPLIED TO OBJECT QUERIES”, filed on 31 Jul. 2015, each of which is hereby incorporated by reference in its entirety for all purposes. See, also, Knowledge Manager Manual, Build a Data Model, Splunk Enterprise 6.1.3 pp. 150-204 (Aug. 25, 2014).

A data model can also include reports. One or more report formats can be associated with a particular data model and be made available to run against the data model. A user can

use child objects to design reports with object datasets that already have extraneous data pre-filtered out. In an embodiment, the data intake and query system **108** provides the user with the ability to produce reports (e.g., a table, chart, visualization, etc.) without having to enter SPL, SQL, or other query language terms into a search screen. Data models are used as the basis for the search feature.

Data models may be selected in a report generation interface. The report generator supports drag-and-drop organization of fields to be summarized in a report. When a model is selected, the fields with available extraction rules are made available for use in the report. The user may refine and/or filter search results to produce more precise reports. The user may select some fields for organizing the report and select other fields for providing detail according to the report organization. For example, “region” and “salesperson” are fields used for organizing the report and sales data can be summarized (subtotaled and totaled) within this organization. The report generator allows the user to specify one or more fields within events and apply statistical analysis on values extracted from the specified one or more fields. The report generator may aggregate search results across sets of events and generate statistics based on aggregated search results. Building reports using the report generation interface is further explained in U.S. patent application Ser. No. 14/503,335, entitled “GENERATING REPORTS FROM UNSTRUCTURED DATA”, filed on 30 Sep. 2014, and which is hereby incorporated by reference in its entirety for all purposes, and in Pivot Manual, Splunk Enterprise 6.1.3 (Aug. 4, 2014). Data visualizations also can be generated in a variety of formats, by reference to the data model. Reports, data visualizations, and data model objects can be saved and associated with the data model for future use. The data model object may be used to perform searches of other data.

FIGS. **12**, **13**, and **7A-7D** illustrate a series of user interface screens where a user may select report generation options using data models. The report generation process may be driven by a predefined data model object, such as a data model object defined and/or saved via a reporting application or a data model object obtained from another source. A user can load a saved data model object using a report editor. For example, the initial search query and fields used to drive the report editor may be obtained from a data model object. The data model object that is used to drive a report generation process may define a search and a set of fields. Upon loading of the data model object, the report generation process may enable a user to use the fields (e.g., the fields defined by the data model object) to define criteria for a report (e.g., filters, split rows/columns, aggregates, etc.) and the search may be used to identify events (e.g., to identify events responsive to the search) used to generate the report. That is, for example, if a data model object is selected to drive a report editor, the graphical user interface of the report editor may enable a user to define reporting criteria for the report using the fields associated with the selected data model object, and the events used to generate the report may be constrained to the events that match, or otherwise satisfy, the search constraints of the selected data model object.

The selection of a data model object for use in driving a report generation may be facilitated by a data model object selection interface. FIG. **12** illustrates an example interactive data model selection graphical user interface **1200** of a report editor that displays a listing of available data models **1201**. The user may select one of the data models **1202**.

FIG. **13** illustrates an example data model object selection graphical user interface **1300** that displays available data

objects **1301** for the selected data object model **1202**. The user may select one of the displayed data model objects **1302** for use in driving the report generation process.

Once a data model object is selected by the user, a user interface screen **700** shown in FIG. 7A may display an interactive listing of automatic field identification options **701** based on the selected data model object. For example, a user may select one of the three illustrated options (e.g., the “All Fields” option **702**, the “Selected Fields” option **703**, or the “Coverage” option (e.g., fields with at least a specified % of coverage) **704**). If the user selects the “All Fields” option **702**, all of the fields identified from the events that were returned in response to an initial search query may be selected. That is, for example, all of the fields of the identified data model object fields may be selected. If the user selects the “Selected Fields” option **703**, only the fields from the fields of the identified data model object fields that are selected by the user may be used. If the user selects the “Coverage” option **704**, only the fields of the identified data model object fields meeting a specified coverage criteria may be selected. A percent coverage may refer to the percentage of events returned by the initial search query that a given field appears in. Thus, for example, if an object dataset includes 10,000 events returned in response to an initial search query, and the “avg\_age” field appears in 854 of those 10,000 events, then the “avg\_age” field would have a coverage of 8.54% for that object dataset. If, for example, the user selects the “Coverage” option and specifies a coverage value of 2%, only fields having a coverage value equal to or greater than 2% may be selected. The number of fields corresponding to each selectable option may be displayed in association with each option. For example, “97” displayed next to the “All Fields” option **702** indicates that 97 fields will be selected if the “All Fields” option is selected. The “3” displayed next to the “Selected Fields” option **703** indicates that 3 of the 97 fields will be selected if the “Selected Fields” option is selected. The “49” displayed next to the “Coverage” option **704** indicates that 49 of the 97 fields (e.g., the 49 fields having a coverage of 2% or greater) will be selected if the “Coverage” option is selected. The number of fields corresponding to the “Coverage” option may be dynamically updated based on the specified percent of coverage.

FIG. 7B illustrates an example graphical user interface screen (also called the pivot interface) **705** displaying the reporting application’s “Report Editor” page. The screen may display interactive elements for defining various elements of a report. For example, the page includes a “Filters” element **706**, a “Split Rows” element **707**, a “Split Columns” element **708**, and a “Column Values” element **709**. The page may include a list of search results **711**. In this example, the Split Rows element **707** is expanded, revealing a listing of fields **710** that can be used to define additional criteria (e.g., reporting criteria). The listing of fields **710** may correspond to the selected fields (attributes). That is, the listing of fields **710** may list only the fields previously selected, either automatically and/or manually by a user. FIG. 7C illustrates a formatting dialogue **712** that may be displayed upon selecting a field from the listing of fields **710**. The dialogue can be used to format the display of the results of the selection (e.g., label the column to be displayed as “component”).

FIG. 7D illustrates an example graphical user interface screen **705** including a table of results **713** based on the selected criteria including splitting the rows by the “component” field. A column **714** having an associated count for each component listed in the table may be displayed that

indicates an aggregate count of the number of times that the particular field-value pair (e.g., the value in a row) occurs in the set of events responsive to the initial search query.

FIG. 14 illustrates an example graphical user interface screen **1400** that allows the user to filter search results and to perform statistical analysis on values extracted from specific fields in the set of events. In this example, the top ten product names ranked by price are selected as a filter **1401** that causes the display of the ten most popular products sorted by price. Each row is displayed by product name and price **1402**. This results in each product displayed in a column labeled “product name” along with an associated price in a column labeled “price” **1406**. Statistical analysis of other fields in the events associated with the ten most popular products have been specified as column values **1403**. A count of the number of successful purchases for each product is displayed in column **1404**. This statistics may be produced by filtering the search results by the product name, finding all occurrences of a successful purchase in a field within the events and generating a total of the number of occurrences. A sum of the total sales is displayed in column **1405**, which is a result of the multiplication of the price and the number of successful purchases for each product.

The reporting application allows the user to create graphical visualizations of the statistics generated for a report. For example, FIG. 15 illustrates an example graphical user interface **1500** that displays a set of components and associated statistics **1501**. The reporting application allows the user to select a visualization of the statistics in a graph (e.g., bar chart, scatter plot, area chart, line chart, pie chart, radial gauge, marker gauge, filler gauge, etc.). FIG. 16 illustrates an example of a bar chart visualization **1600** of an aspect of the statistical data **1501**. FIG. 17 illustrates a scatter plot visualization **1700** of an aspect of the statistical data **1501**.

## 2.10. Acceleration Technique

The above-described system provides significant flexibility by enabling a user to analyze massive quantities of minimally processed data “on the fly” at search time instead of storing pre-specified portions of the data in a database at ingestion time. This flexibility enables a user to see valuable insights, correlate data, and perform subsequent queries to examine interesting aspects of the data that may not have been apparent at ingestion time.

However, performing extraction and analysis operations at search time can involve a large amount of data and require a large number of computational operations, which can cause delays in processing the queries. Advantageously, SPLUNK® ENTERPRISE system employs a number of unique acceleration techniques that have been developed to speed up analysis operations performed at search time. These techniques include: (1) performing search operations in parallel across multiple indexers; (2) using a keyword index; (3) using a high performance analytics store; and (4) accelerating the process of generating reports. These novel techniques are described in more detail below.

### 2.10.1. Aggregation Technique

To facilitate faster query processing, a query can be structured such that multiple indexers perform the query in parallel, while aggregation of search results from the multiple indexers is performed locally at the search head. For example, FIG. 8 illustrates how a search query **802** received from a client at a search head **210** can split into two phases,

including: (1) subtasks **804** (e.g., data retrieval or simple filtering) that may be performed in parallel by indexers **206** for execution, and (2) a search results aggregation operation **806** to be executed by the search head when the results are ultimately collected from the indexers.

During operation, upon receiving search query **802**, a search head **210** determines that a portion of the operations involved with the search query may be performed locally by the search head. The search head modifies search query **802** by substituting “stats” (create aggregate statistics over results sets received from the indexers at the search head) with “prestats” (create statistics by the indexer from local results set) to produce search query **804**, and then distributes search query **804** to distributed indexers, which are also referred to as “search peers.” Note that search queries may generally specify search criteria or operations to be performed on events that meet the search criteria. Search queries may also specify field names, as well as search criteria for the values in the fields or operations to be performed on the values in the fields. Moreover, the search head may distribute the full search query to the search peers as illustrated in FIG. 4, or may alternatively distribute a modified version (e.g., a more restricted version) of the search query to the search peers. In this example, the indexers are responsible for producing the results and sending them to the search head. After the indexers return the results to the search head, the search head aggregates the received results **806** to form a single search result set. By executing the query in this manner, the system effectively distributes the computational operations across the indexers while minimizing data transfers.

#### 2.10.2. Keyword Index

As described above with reference to the flow charts in FIG. 3 and FIG. 4, data intake and query system **108** can construct and maintain one or more keyword indices to quickly identify events containing specific keywords. This technique can greatly speed up the processing of queries involving specific keywords. As mentioned above, to build a keyword index, an indexer first identifies a set of keywords. Then, the indexer includes the identified keywords in an index, which associates each stored keyword with references to events containing that keyword, or to locations within events where that keyword is located. When an indexer subsequently receives a keyword-based query, the indexer can access the keyword index to quickly identify events containing the keyword.

#### 2.10.3. High Performance Analytics Store

To speed up certain types of queries, some embodiments of system **108** create a high performance analytics store, which is referred to as a “summarization table,” that contains entries for specific field-value pairs. Each of these entries keeps track of instances of a specific value in a specific field in the event data and includes references to events containing the specific value in the specific field. For example, an example entry in a summarization table can keep track of occurrences of the value “94107” in a “ZIP code” field of a set of events and the entry includes references to all of the events that contain the value “94107” in the ZIP code field. This optimization technique enables the system to quickly process queries that seek to determine how many events have a particular value for a particular field. To this end, the system can examine the entry in the summarization table to count instances of the specific value in the field without

having to go through the individual events or perform data extractions at search time. Also, if the system needs to process all events that have a specific field-value combination, the system can use the references in the summarization table entry to directly access the events to extract further information without having to search all of the events to find the specific field-value combination at search time.

In some embodiments, the system maintains a separate summarization table for each of the above-described time-specific buckets that stores events for a specific time range. A bucket-specific summarization table includes entries for specific field-value combinations that occur in events in the specific bucket. Alternatively, the system can maintain a separate summarization table for each indexer. The indexer-specific summarization table includes entries for the events in a data store that are managed by the specific indexer. Indexer-specific summarization tables may also be bucket-specific.

The summarization table can be populated by running a periodic query that scans a set of events to find instances of a specific field-value combination, or alternatively instances of all field-value combinations for a specific field. A periodic query can be initiated by a user, or can be scheduled to occur automatically at specific time intervals. A periodic query can also be automatically launched in response to a query that asks for a specific field-value combination.

In some cases, when the summarization tables may not cover all of the events that are relevant to a query, the system can use the summarization tables to obtain partial results for the events that are covered by summarization tables, but may also have to search through other events that are not covered by the summarization tables to produce additional results. These additional results can then be combined with the partial results to produce a final set of results for the query. The summarization table and associated techniques are described in more detail in U.S. Pat. No. 8,682,925, entitled “DISTRIBUTED HIGH PERFORMANCE ANALYTICS STORE”, issued on 25 Mar. 2014, U.S. patent application Ser. No. 14/170,159, entitled “SUPPLEMENTING A HIGH PERFORMANCE ANALYTICS STORE WITH EVALUATION OF INDIVIDUAL EVENTS TO RESPOND TO AN EVENT QUERY”, filed on 31 Jan. 2014, and U.S. patent application Ser. No. 14/815,973, entitled “STORAGE MEDIUM AND CONTROL DEVICE”, filed on 21 Feb. 2014, each of which is hereby incorporated by reference in its entirety.

#### 2.10.4. Accelerating Report Generation

In some embodiments, a data server system such as the SPLUNK® ENTERPRISE system can accelerate the process of periodically generating updated reports based on query results. To accelerate this process, a summarization engine automatically examines the query to determine whether generation of updated reports can be accelerated by creating intermediate summaries. If reports can be accelerated, the summarization engine periodically generates a summary covering data obtained during a latest non-overlapping time period. For example, where the query seeks events meeting a specified criteria, a summary for the time period includes only events within the time period that meet the specified criteria. Similarly, if the query seeks statistics calculated from the events, such as the number of events that match the specified criteria, then the summary for the time period includes the number of events in the period that match the specified criteria.

In addition to the creation of the summaries, the summarization engine schedules the periodic updating of the report associated with the query. During each scheduled report update, the query engine determines whether intermediate summaries have been generated covering portions of the time period covered by the report update. If so, then the report is generated based on the information contained in the summaries. Also, if additional event data has been received and has not yet been summarized, and is required to generate the complete report, the query can be run on this additional event data. Then, the results returned by this query on the additional event data, along with the partial results obtained from the intermediate summaries, can be combined to generate the updated report. This process is repeated each time the report is updated. Alternatively, if the system stores events in buckets covering specific time ranges, then the summaries can be generated on a bucket-by-bucket basis. Note that producing intermediate summaries can save the work involved in re-running the query for previous time periods, so advantageously only the newer event data needs to be processed while generating an updated report. These report acceleration techniques are described in more detail in U.S. Pat. No. 8,589,403, entitled "COMPRESSED JOURNALING IN EVENT TRACKING FILES FOR METADATA RECOVERY AND REPLICATION", issued on 19 Nov. 2013, U.S. Pat. No. 8,412,696, entitled "REAL TIME SEARCHING AND REPORTING", issued on 2 Apr. 2011, and U.S. Pat. Nos. 8,589,375 and 8,589,432, both also entitled "REAL TIME SEARCHING AND REPORTING", both issued on 19 Nov. 2013, each of which is hereby incorporated by reference in its entirety.

### 2.11. Security Features

The SPLUNK® ENTERPRISE platform provides various schemas, dashboards and visualizations that simplify developers' task to create applications with additional capabilities. One such application is the SPLUNK® APP FOR ENTERPRISE SECURITY, which performs monitoring and alerting operations and includes analytics to facilitate identifying both known and unknown security threats based on large volumes of data stored by the SPLUNK® ENTERPRISE system. SPLUNK® APP FOR ENTERPRISE SECURITY provides the security practitioner with visibility into security-relevant threats found in the enterprise infrastructure by capturing, monitoring, and reporting on data from enterprise security devices, systems, and applications. Through the use of SPLUNK® ENTERPRISE searching and reporting capabilities, SPLUNK® APP FOR ENTERPRISE SECURITY provides a top-down and bottom-up view of an organization's security posture.

The SPLUNK® APP FOR ENTERPRISE SECURITY leverages SPLUNK® ENTERPRISE search-time normalization techniques, saved searches, and correlation searches to provide visibility into security-relevant threats and activity and generate notable events for tracking. The App enables the security practitioner to investigate and explore the data to find new or unknown threats that do not follow signature-based patterns.

Conventional Security Information and Event Management (SIEM) systems that lack the infrastructure to effectively store and analyze large volumes of security-related data. Traditional SIEM systems typically use fixed schemas to extract data from pre-defined security-related fields at data ingestion time and storing the extracted data in a relational database. This traditional data extraction process (and associated reduction in data size) that occurs at data

ingestion time inevitably hampers future incident investigations that may need original data to determine the root cause of a security issue, or to detect the onset of an impending security threat.

In contrast, the SPLUNK® APP FOR ENTERPRISE SECURITY system stores large volumes of minimally processed security-related data at ingestion time for later retrieval and analysis at search time when a live security threat is being investigated. To facilitate this data retrieval process, the SPLUNK® APP FOR ENTERPRISE SECURITY provides pre-specified schemas for extracting relevant values from the different types of security-related event data and enables a user to define such schemas.

The SPLUNK® APP FOR ENTERPRISE SECURITY can process many types of security-related information. In general, this security-related information can include any information that can be used to identify security threats. For example, the security-related information can include network-related information, such as IP addresses, domain names, asset identifiers, network traffic volume, uniform resource locator strings, and source addresses. The process of detecting security threats for network-related information is further described in U.S. Pat. No. 8,826,434, entitled "SECURITY THREAT DETECTION BASED ON INDICATIONS IN BIG DATA OF ACCESS TO NEWLY REGISTERED DOMAINS", issued on 2 Sep. 2014, U.S. patent application Ser. No. 13/956,252, entitled "INVESTIGATIVE AND DYNAMIC DETECTION OF POTENTIAL SECURITY-THREAT INDICATORS FROM EVENTS IN BIG DATA", filed on 31 Jul. 2013, U.S. patent application Ser. No. 14/445,018, entitled "GRAPHIC DISPLAY OF SECURITY THREATS BASED ON INDICATIONS OF ACCESS TO NEWLY REGISTERED DOMAINS", filed on 28 Jul. 2014, U.S. patent application Ser. No. 14/445,023, entitled "SECURITY THREAT DETECTION OF NEWLY REGISTERED DOMAINS", filed on 28 Jul. 2014, U.S. patent application Ser. No. 14/815,971, entitled "SECURITY THREAT DETECTION USING DOMAIN NAME ACCESSES", filed on 1 Aug. 2015, and U.S. patent application Ser. No. 14/815,972, entitled "SECURITY THREAT DETECTION USING DOMAIN NAME REGISTRATIONS", filed on 1 Aug. 2015, each of which is hereby incorporated by reference in its entirety for all purposes. Security-related information can also include malware infection data and system configuration information, as well as access control information, such as login/logout information and access failure notifications. The security-related information can originate from various sources within a data center, such as hosts, virtual machines, storage devices and sensors. The security-related information can also originate from various sources in a network, such as routers, switches, email servers, proxy servers, gateways, firewalls and intrusion-detection systems.

During operation, the SPLUNK® APP FOR ENTERPRISE SECURITY facilitates detecting "notable events" that are likely to indicate a security threat. These notable events can be detected in a number of ways: (1) a user can notice a correlation in the data and can manually identify a corresponding group of one or more events as "notable;" or (2) a user can define a "correlation search" specifying criteria for a notable event, and every time one or more events satisfy the criteria, the application can indicate that the one or more events are notable. A user can alternatively select a pre-defined correlation search provided by the application. Note that correlation searches can be run continuously or at regular intervals (e.g., every hour) to search for notable events. Upon detection, notable events can be

stored in a dedicated “notable events index,” which can be subsequently accessed to generate various visualizations containing security-related information. Also, alerts can be generated to notify system operators when important notable events are discovered.

The SPLUNK® APP FOR ENTERPRISE SECURITY provides various visualizations to aid in discovering security threats, such as a “key indicators view” that enables a user to view security metrics, such as counts of different types of notable events. For example, FIG. 9A illustrates an example key indicators view 900 that comprises a dashboard, which can display a value 901, for various security-related metrics, such as malware infections 902. It can also display a change in a metric value 903, which indicates that the number of malware infections increased by 63 during the preceding interval. Key indicators view 900 additionally displays a histogram panel 904 that displays a histogram of notable events organized by urgency values, and a histogram of notable events organized by time intervals. This key indicators view is described in further detail in pending U.S. patent application Ser. No. 13/956,338, entitled “KEY INDICATORS VIEW”, filed on 31 Jul. 2013, and which is hereby incorporated by reference in its entirety for all purposes.

These visualizations can also include an “incident review dashboard” that enables a user to view and act on “notable events.” These notable events can include: (1) a single event of high importance, such as any activity from a known web attacker; or (2) multiple events that collectively warrant review, such as a large number of authentication failures on a host followed by a successful authentication. For example, FIG. 9B illustrates an example incident review dashboard 910 that includes a set of incident attribute fields 911 that, for example, enables a user to specify a time range field 912 for the displayed events. It also includes a timeline 913 that graphically illustrates the number of incidents that occurred in time intervals over the selected time range. It additionally displays an events list 914 that enables a user to view a list of all of the notable events that match the criteria in the incident attributes fields 911. To facilitate identifying patterns among the notable events, each notable event can be associated with an urgency value (e.g., low, medium, high, critical), which is indicated in the incident review dashboard. The urgency value for a detected event can be determined based on the severity of the event and the priority of the system component associated with the event.

## 2.12. Data Center Monitoring

As mentioned above, the SPLUNK® ENTERPRISE platform provides various features that simplify the developers’ task to create various applications. One such application is SPLUNK® APP FOR VMWARE® that provides operational visibility into granular performance metrics, logs, tasks and events, and topology from hosts, virtual machines and virtual centers. It empowers administrators with an accurate real-time picture of the health of the environment, proactively identifying performance and capacity bottlenecks.

Conventional data-center-monitoring systems lack the infrastructure to effectively store and analyze large volumes of machine-generated data, such as performance information and log data obtained from the data center. In conventional data-center-monitoring systems, machine-generated data is typically pre-processed prior to being stored, for example, by extracting pre-specified data items and storing them in a database to facilitate subsequent retrieval and analysis at

search time. However, the rest of the data is not saved and discarded during pre-processing.

In contrast, the SPLUNK® APP FOR VMWARE® stores large volumes of minimally processed machine data, such as performance information and log data, at ingestion time for later retrieval and analysis at search time when a live performance issue is being investigated. In addition to data obtained from various log files, this performance-related information can include values for performance metrics obtained through an application programming interface (API) provided as part of the vSphere Hypervisor™ system distributed by VMware, Inc. of Palo Alto, Calif. For example, these performance metrics can include: (1) CPU-related performance metrics; (2) disk-related performance metrics; (3) memory-related performance metrics; (4) network-related performance metrics; (5) energy-usage statistics; (6) data-traffic-related performance metrics; (7) overall system availability performance metrics; (8) cluster-related performance metrics; and (9) virtual machine performance statistics. Such performance metrics are described in U.S. patent application Ser. No. 14/167,316, entitled “CORRELATION FOR USER-SELECTED TIME RANGES OF VALUES FOR PERFORMANCE METRICS OF COMPONENTS IN AN INFORMATION-TECHNOLOGY ENVIRONMENT WITH LOG DATA FROM THAT INFORMATION-TECHNOLOGY ENVIRONMENT”, filed on 29 Jan. 2014, and which is hereby incorporated by reference in its entirety for all purposes.

To facilitate retrieving information of interest from performance data and log files, the SPLUNK® APP FOR VMWARE® provides pre-specified schemas for extracting relevant values from different types of performance-related event data, and also enables a user to define such schemas.

The SPLUNK® APP FOR VMWARE® additionally provides various visualizations to facilitate detecting and diagnosing the root cause of performance problems. For example, one such visualization is a “proactive monitoring tree” that enables a user to easily view and understand relationships among various factors that affect the performance of a hierarchically structured computing system. This proactive monitoring tree enables a user to easily navigate the hierarchy by selectively expanding nodes representing various entities (e.g., virtual centers or computing clusters) to view performance information for lower-level nodes associated with lower-level entities (e.g., virtual machines or host systems). Example node-expansion operations are illustrated in FIG. 9C, wherein nodes 933 and 934 are selectively expanded. Note that nodes 931-939 can be displayed using different patterns or colors to represent different performance states, such as a critical state, a warning state, a normal state or an unknown/offline state. The ease of navigation provided by selective expansion in combination with the associated performance-state information enables a user to quickly diagnose the root cause of a performance problem. The proactive monitoring tree is described in further detail in U.S. patent application Ser. No. 14/253,490, entitled “PROACTIVE MONITORING TREE WITH SEVERITY STATE SORTING”, filed on 15 Apr. 2014, and U.S. patent application Ser. No. 14/812,948, also entitled “PROACTIVE MONITORING TREE WITH SEVERITY STATE SORTING”, filed on 29 Jul. 2015, each of which is hereby incorporated by reference in its entirety for all purposes.

The SPLUNK® APP FOR VMWARE® also provides a user interface that enables a user to select a specific time range and then view heterogeneous data comprising events, log data, and associated performance metrics for the selected

time range. For example, the screen illustrated in FIG. 9D displays a listing of recent “tasks and events” and a listing of recent “log entries” for a selected time range above a performance-metric graph for “average CPU core utilization” for the selected time range. Note that a user is able to operate pull-down menus 942 to selectively display different performance metric graphs for the selected time range. This enables the user to correlate trends in the performance-metric graph with corresponding event and log data to quickly determine the root cause of a performance problem. This user interface is described in more detail in U.S. patent application Ser. No. 14/167,316, entitled “CORRELATION FOR USER-SELECTED TIME RANGES OF VALUES FOR PERFORMANCE METRICS OF COMPONENTS IN AN INFORMATION-TECHNOLOGY ENVIRONMENT WITH LOG DATA FROM THAT INFORMATION-TECHNOLOGY ENVIRONMENT”, filed on 29 Jan. 2014, and which is hereby incorporated by reference in its entirety for all purposes.

### 2.13. Cloud-Based System Overview

The example data intake and query system 108 described in reference to FIG. 2 comprises several system components, including one or more forwarders, indexers, and search heads. In some environments, a user of a data intake and query system 108 may install and configure, on computing devices owned and operated by the user, one or more software applications that implement some or all of these system components. For example, a user may install a software application on server computers owned by the user and configure each server to operate as one or more of a forwarder, an indexer, a search head, etc. This arrangement generally may be referred to as an “on-premises” solution. That is, the system 108 is installed and operates on computing devices directly controlled by the user of the system. Some users may prefer an on-premises solution because it may provide a greater level of control over the configuration of certain aspects of the system (e.g., security, privacy, standards, controls, etc.). However, other users may instead prefer an arrangement in which the user is not directly responsible for providing and managing the computing devices upon which various components of system 108 operate.

In one embodiment, to provide an alternative to an entirely on-premises environment for system 108, one or more of the components of a data intake and query system instead may be provided as a cloud-based service. In this context, a cloud-based service refers to a service hosted by one or more computing resources that are accessible to end users over a network, for example, by using a web browser or other application on a client device to interface with the remote computing resources. For example, a service provider may provide a cloud-based data intake and query system by managing computing resources configured to implement various aspects of the system (e.g., forwarders, indexers, search heads, etc.) and by providing access to the system to end users via a network. Typically, a user may pay a subscription or other fee to use such a service. Each subscribing user of the cloud-based service may be provided with an account that enables the user to configure a customized cloud-based system based on the user’s preferences.

FIG. 10 illustrates a block diagram of an example cloud-based data intake and query system. Similar to the system of FIG. 2, the networked computer system 1000 includes input data sources 202 and forwarders 204. These input data sources and forwarders may be in a subscriber’s private

computing environment. Alternatively, they might be directly managed by the service provider as part of the cloud service. In the example system 1000, one or more forwarders 204 and client devices 1002 are coupled to a cloud-based data intake and query system 1006 via one or more networks 1004. Network 1004 broadly represents one or more LANs, WANs, cellular networks, intranetworks, internetworks, etc., using any of wired, wireless, terrestrial microwave, satellite links, etc., and may include the public Internet, and is used by client devices 1002 and forwarders 204 to access the system 1006. Similar to the system of 108, each of the forwarders 204 may be configured to receive data from an input source and to forward the data to other components of the system 1006 for further processing.

In an embodiment, a cloud-based data intake and query system 1006 may comprise a plurality of system instances 1008. In general, each system instance 1008 may include one or more computing resources managed by a provider of the cloud-based system 1006 made available to a particular subscriber. The computing resources comprising a system instance 1008 may, for example, include one or more servers or other devices configured to implement one or more forwarders, indexers, search heads, and other components of a data intake and query system, similar to system 108. As indicated above, a subscriber may use a web browser or other application of a client device 1002 to access a web portal or other interface that enables the subscriber to configure an instance 1008.

Providing a data intake and query system as described in reference to system 108 as a cloud-based service presents a number of challenges. Each of the components of a system 108 (e.g., forwarders, indexers and search heads) may at times refer to various configuration files stored locally at each component. These configuration files typically may involve some level of user configuration to accommodate particular types of data a user desires to analyze and to account for other user preferences. However, in a cloud-based service context, users typically may not have direct access to the underlying computing resources implementing the various system components (e.g., the computing resources comprising each system instance 1008) and may desire to make such configurations indirectly, for example, using one or more web-based interfaces. Thus, the techniques and systems described herein for providing user interfaces that enable a user to configure source type definitions are applicable to both on-premises and cloud-based service contexts, or some combination thereof (e.g., a hybrid system where both an on-premises environment such as SPLUNK® ENTERPRISE and a cloud-based environment such as SPLUNK CLOUD™ are centrally visible).

### 2.14. Searching Externally Archived Data

FIG. 11 shows a block diagram of an example of a data intake and query system 108 that provides transparent search facilities for data systems that are external to the data intake and query system. Such facilities are available in the HUNK® system provided by Splunk Inc. of San Francisco, Calif. HUNK® represents an analytics platform that enables business and IT teams to rapidly explore, analyze, and visualize data in Hadoop and NoSQL data stores.

The search head 210 of the data intake and query system receives search requests from one or more client devices 1104 over network connections 1120. As discussed above, the data intake and query system 108 may reside in an enterprise location, in the cloud, etc. FIG. 11 illustrates that multiple client devices 1104a, 1104b, . . . , 1104n may

communicate with the data intake and query system **108**. The client devices **1104** may communicate with the data intake and query system using a variety of connections. For example, one client device in FIG. **11** is illustrated as communicating over an Internet (Web) protocol, another client device is illustrated as communicating via a command line interface, and another client device is illustrated as communicating via a system developer kit (SDK).

The search head **210** analyzes the received search request to identify request parameters. If a search request received from one of the client devices **1104** references an index maintained by the data intake and query system, then the search head **210** connects to one or more indexers **206** of the data intake and query system for the index referenced in the request parameters. That is, if the request parameters of the search request reference an index, then the search head accesses the data in the index via the indexer. The data intake and query system **108** may include one or more indexers **206**, depending on system access resources and requirements. As described further below, the indexers **206** retrieve data from their respective local data stores **208** as specified in the search request. The indexers and their respective data stores can comprise one or more storage devices and typically reside on the same system, though they may be connected via a local network connection.

If the request parameters of the received search request reference an external data collection, which is not accessible to the indexers **206** or under the management of the data intake and query system, then the search head **210** can access the external data collection through an External Result Provider (ERP) process **1110**. An external data collection may be referred to as a “virtual index” (plural, “virtual indices”). An ERP process provides an interface through which the search head **210** may access virtual indices.

Thus, a search reference to an index of the system relates to a locally stored and managed data collection. In contrast, a search reference to a virtual index relates to an externally stored and managed data collection, which the search head may access through one or more ERP processes **1110**, **1112**. FIG. **11** shows two ERP processes **1110**, **1112** that connect to respective remote (external) virtual indices, which are indicated as a Hadoop or another system **1114** (e.g., Amazon S3, Amazon EMR, other Hadoop Compatible File Systems (HCFS), etc.) and a relational database management system (RDBMS) **1116**. Other virtual indices may include other file organizations and protocols, such as Structured Query Language (SQL) and the like. The ellipses between the ERP processes **1110**, **1112** indicate optional additional ERP processes of the data intake and query system **108**. An ERP process may be a computer process that is initiated or spawned by the search head **210** and is executed by the search data intake and query system **108**. Alternatively or additionally, an ERP process may be a process spawned by the search head **210** on the same or different host system as the search head **210** resides.

The search head **210** may spawn a single ERP process in response to multiple virtual indices referenced in a search request, or the search head may spawn different ERP processes for different virtual indices. Generally, virtual indices that share common data configurations or protocols may share ERP processes. For example, all search query references to a Hadoop file system may be processed by the same ERP process, if the ERP process is suitably configured. Likewise, all search query references to an SQL database may be processed by the same ERP process. In addition, the search head may provide a common ERP process for com-

mon external data source types (e.g., a common vendor may utilize a common ERP process, even if the vendor includes different data storage system types, such as Hadoop and SQL). Common indexing schemes also may be handled by common ERP processes, such as flat text files or Weblog files.

The search head **210** determines the number of ERP processes to be initiated via the use of configuration parameters that are included in a search request message. Generally, there is a one-to-many relationship between an external results provider “family” and ERP processes. There is also a one-to-many relationship between an ERP process and corresponding virtual indices that are referred to in a search request. For example, using RDBMS, assume two independent instances of such a system by one vendor, such as one RDBMS for production and another RDBMS used for development. In such a situation, it is likely preferable (but optional) to use two ERP processes to maintain the independent operation as between production and development data. Both of the ERPs, however, will belong to the same family, because the two RDBMS system types are from the same vendor.

The ERP processes **1110**, **1112** receive a search request from the search head **210**. The search head may optimize the received search request for execution at the respective external virtual index. Alternatively, the ERP process may receive a search request as a result of analysis performed by the search head or by a different system process. The ERP processes **1110**, **1112** can communicate with the search head **210** via conventional input/output routines (e.g., standard in/standard out, etc.). In this way, the ERP process receives the search request from a client device such that the search request may be efficiently executed at the corresponding external virtual index.

The ERP processes **1110**, **1112** may be implemented as a process of the data intake and query system. Each ERP process may be provided by the data intake and query system, or may be provided by process or application providers who are independent of the data intake and query system. Each respective ERP process may include an interface application installed at a computer of the external result provider that ensures proper communication between the search support system and the external result provider. The ERP processes **1110**, **1112** generate appropriate search requests in the protocol and syntax of the respective virtual indices **1114**, **1116**, each of which corresponds to the search request received by the search head **210**. Upon receiving search results from their corresponding virtual indices, the respective ERP process passes the result to the search head **210**, which may return or display the results or a processed set of results based on the returned results to the respective client device.

Client devices **1104** may communicate with the data intake and query system **108** through a network interface **1120**, e.g., one or more LANs, WANs, cellular networks, intranetworks, and/or internetworks using any of wired, wireless, terrestrial microwave, satellite links, etc., and may include the public Internet.

The analytics platform utilizing the External Result Provider process described in more detail in U.S. Pat. No. 8,738,629, entitled “EXTERNAL RESULT PROVIDED PROCESS FOR RETRIEVING DATA STORED USING A DIFFERENT CONFIGURATION OR PROTOCOL”, issued on 27 May 2014, U.S. Pat. No. 8,738,587, entitled “PROCESSING A SYSTEM SEARCH REQUEST BY RETRIEVING RESULTS FROM BOTH A NATIVE INDEX AND A VIRTUAL INDEX”, issued on 25 Jul. 2013,

U.S. patent application Ser. No. 14/266,832, entitled “PROCESSING A SYSTEM SEARCH REQUEST ACROSS DISPARATE DATA COLLECTION SYSTEMS”, filed on 1 May 2014, and U.S. patent application Ser. No. 14/449,144, entitled “PROCESSING A SYSTEM SEARCH REQUEST INCLUDING EXTERNAL DATA SOURCES”, filed on 31 Jul. 2014, each of which is hereby incorporated by reference in its entirety for all purposes.

#### 2.14.1. ERP Process Features

The ERP processes described above may include two operation modes: a streaming mode and a reporting mode. The ERP processes can operate in streaming mode only, in reporting mode only, or in both modes simultaneously. Operating in both modes simultaneously is referred to as mixed mode operation. In a mixed mode operation, the ERP at some point can stop providing the search head with streaming results and only provide reporting results thereafter, or the search head at some point may start ignoring streaming results it has been using and only use reporting results thereafter.

The streaming mode returns search results in real time, with minimal processing, in response to the search request. The reporting mode provides results of a search request with processing of the search results prior to providing them to the requesting search head, which in turn provides results to the requesting client device. ERP operation with such multiple modes provides greater performance flexibility with regard to report time, search latency, and resource utilization.

In a mixed mode operation, both streaming mode and reporting mode are operating simultaneously. The streaming mode results (e.g., the raw data obtained from the external data source) are provided to the search head, which can then process the results data (e.g., break the raw data into events, timestamp it, filter it, etc.) and integrate the results data with the results data from other external data sources, and/or from data stores of the search head. The search head performs such processing and can immediately start returning interim (streaming mode) results to the user at the requesting client device; simultaneously, the search head is waiting for the ERP process to process the data it is retrieving from the external data source as a result of the concurrently executing reporting mode.

In some instances, the ERP process initially operates in a mixed mode, such that the streaming mode operates to enable the ERP quickly to return interim results (e.g., some of the raw or unprocessed data necessary to respond to a search request) to the search head, enabling the search head to process the interim results and begin providing to the client or search requester interim results that are responsive to the query. Meanwhile, in this mixed mode, the ERP also operates concurrently in reporting mode, processing portions of raw data in a manner responsive to the search query. Upon determining that it has results from the reporting mode available to return to the search head, the ERP may halt processing in the mixed mode at that time (or some later time) by stopping the return of data in streaming mode to the search head and switching to reporting mode only. The ERP at this point starts sending interim results in reporting mode to the search head, which in turn may then present this processed data responsive to the search request to the client or search requester. Typically the search head switches from using results from the ERP’s streaming mode of operation to results from the ERP’s reporting mode of operation when the higher bandwidth results from the reporting mode out-

strip the amount of data processed by the search head in the streaming mode of ERP operation.

A reporting mode may have a higher bandwidth because the ERP does not have to spend time transferring data to the search head for processing all the raw data. In addition, the ERP may optionally direct another processor to do the processing.

The streaming mode of operation does not need to be stopped to gain the higher bandwidth benefits of a reporting mode; the search head could simply stop using the streaming mode results—and start using the reporting mode results—when the bandwidth of the reporting mode has caught up with or exceeded the amount of bandwidth provided by the streaming mode. Thus, a variety of triggers and ways to accomplish a search head’s switch from using streaming mode results to using reporting mode results may be appreciated by one skilled in the art.

The reporting mode can involve the ERP process (or an external system) performing event breaking, time stamping, filtering of events to match the search query request, and calculating statistics on the results. The user can request particular types of data, such as if the search query itself involves types of events, or the search request may ask for statistics on data, such as on events that meet the search request. In either case, the search head understands the query language used in the received query request, which may be a proprietary language. One exemplary query language is Splunk Processing Language (SPL) developed by the assignee of the application, Splunk Inc. The search head typically understands how to use that language to obtain data from the indexers, which store data in a format used by the SPLUNK® Enterprise system.

The ERP processes support the search head, as the search head is not ordinarily configured to understand the format in which data is stored in external data sources such as Hadoop or SQL data systems. Rather, the ERP process performs that translation from the query submitted in the search support system’s native format (e.g., SPL if SPLUNK® ENTERPRISE is used as the search support system) to a search query request format that will be accepted by the corresponding external data system. The external data system typically stores data in a different format from that of the search support system’s native index format, and it utilizes a different query language (e.g., SQL or MapReduce, rather than SPL or the like).

As noted, the ERP process can operate in the streaming mode alone. After the ERP process has performed the translation of the query request and received raw results from the streaming mode, the search head can integrate the returned data with any data obtained from local data sources (e.g., native to the search support system), other external data sources, and other ERP processes (if such operations were required to satisfy the terms of the search query). An advantage of mixed mode operation is that, in addition to streaming mode, the ERP process is also executing concurrently in reporting mode. Thus, the ERP process (rather than the search head) is processing query results (e.g., performing event breaking, timestamping, filtering, possibly calculating statistics if required to be responsive to the search query request, etc.). It should be apparent to those skilled in the art that additional time is needed for the ERP process to perform the processing in such a configuration. Therefore, the streaming mode will allow the search head to start returning interim results to the user at the client device before the ERP process can complete sufficient processing to start returning any search results. The switchover between streaming and reporting mode happens when the ERP process determines



that the switchover is appropriate, such as when the ERP process determines it can begin returning meaningful results from its reporting mode.

The operation described above illustrates the source of operational latency: streaming mode has low latency (immediate results) and usually has relatively low bandwidth (fewer results can be returned per unit of time). In contrast, the concurrently running reporting mode has relatively high latency (it has to perform a lot more processing before returning any results) and usually has relatively high bandwidth (more results can be processed per unit of time). For example, when the ERP process does begin returning report results, it returns more processed results than in the streaming mode, because, e.g., statistics only need to be calculated to be responsive to the search request. That is, the ERP process doesn't have to take time to first return raw data to the search head. As noted, the ERP process could be configured to operate in streaming mode alone and return just the raw data for the search head to process in a way that is responsive to the search request. Alternatively, the ERP process can be configured to operate in the reporting mode only. Also, the ERP process can be configured to operate in streaming mode and reporting mode concurrently, as described, with the ERP process stopping the transmission of streaming results to the search head when the concurrently running reporting mode has caught up and started providing results. The reporting mode does not require the processing of all raw data that is responsive to the search query request before the ERP process starts returning results; rather, the reporting mode usually performs processing of chunks of events and returns the processing results to the search head for each chunk.

For example, an ERP process can be configured to merely return the contents of a search result file verbatim, with little or no processing of results. That way, the search head performs all processing (such as parsing byte streams into events, filtering, etc.). The ERP process can be configured to perform additional intelligence, such as analyzing the search request and handling all the computation that a native search indexer process would otherwise perform. In this way, the configured ERP process provides greater flexibility in features while operating according to desired preferences, such as response latency and resource requirements.

### 2.15. IT Service Monitoring

As previously mentioned, the SPLUNK® ENTERPRISE platform provides various schemas, dashboards and visualizations that make it easy for developers to create applications to provide additional capabilities. One such application is SPLUNK® IT SERVICE INTELLIGENCE™, which performs monitoring and alerting operations. It also includes analytics to help an analyst diagnose the root cause of performance problems based on large volumes of data stored by the SPLUNK® ENTERPRISE system as correlated to the various services an IT organization provides (a service-centric view). This differs significantly from conventional IT monitoring systems that lack the infrastructure to effectively store and analyze large volumes of service-related event data. Traditional service monitoring systems typically use fixed schemas to extract data from pre-defined fields at data ingestion time, wherein the extracted data is typically stored in a relational database. This data extraction process and associated reduction in data content that occurs at data ingestion time inevitably hampers future investigations, when all of the original data may be needed to determine the root cause of or contributing factors to a service issue.

In contrast, a SPLUNK® IT SERVICE INTELLIGENCE™ system stores large volumes of minimally-processed service-related data at ingestion time for later retrieval and analysis at search time, to perform regular monitoring, or to investigate a service issue. To facilitate this data retrieval process, SPLUNK® IT SERVICE INTELLIGENCE™ enables a user to define an IT operations infrastructure from the perspective of the services it provides. In this service-centric approach, a service such as corporate e-mail may be defined in terms of the entities employed to provide the service, such as host machines and network devices. Each entity is defined to include information for identifying all of the event data that pertains to the entity, whether produced by the entity itself or by another machine, and considering the many various ways the entity may be identified in raw machine data (such as by a URL, an IP address, or machine name). The service and entity definitions can organize event data around a service so that all of the event data pertaining to that service can be easily identified. This capability provides a foundation for the implementation of Key Performance Indicators.

One or more Key Performance Indicators (KPI's) are defined for a service within the SPLUNK® IT SERVICE INTELLIGENCE™ application. Each KPI measures an aspect of service performance at a point in time or over a period of time (aspect KPI's). Each KPI is defined by a search query that derives a KPI value from the machine data of events associated with the entities that provide the service. Information in the entity definitions may be used to identify the appropriate events at the time a KPI is defined or whenever a KPI value is being determined. The KPI values derived over time may be stored to build a valuable repository of current and historical performance information for the service, and the repository, itself, may be subject to search query processing. Aggregate KPIs may be defined to provide a measure of service performance calculated from a set of service aspect KPI values; this aggregate may even be taken across defined timeframes and/or across multiple services. A particular service may have an aggregate KPI derived from substantially all of the aspect KPI's of the service to indicate an overall health score for the service.

SPLUNK® IT SERVICE INTELLIGENCE™ facilitates the production of meaningful aggregate KPI's through a system of KPI thresholds and state values. Different KPI definitions may produce values in different ranges, and so the same value may mean something very different from one KPI definition to another. To address this, SPLUNK® IT SERVICE INTELLIGENCE™ implements a translation of individual KPI values to a common domain of "state" values. For example, a KPI range of values may be 1-100, or 50-275, while values in the state domain may be 'critical,' 'warning,' 'normal,' and 'informational' . . . . Thresholds associated with a particular KPI definition determine ranges of values for that KPI that correspond to the various state values. In one case, KPI values **95-100** may be set to correspond to 'critical' in the state domain. KPI values from disparate KPI's can be processed uniformly once they are translated into the common state values using the thresholds. For example, "normal 80% of the time" can be applied across various KPI's. To provide meaningful aggregate KPI's, a weighting value can be assigned to each KPI so that its influence on the calculated aggregate KPI value is increased or decreased relative to the other KPI's.

One service in an IT environment often impacts, or is impacted by, another service. SPLUNK® IT SERVICE INTELLIGENCE™ can reflect these dependencies. For example, a dependency relationship between a corporate

e-mail service and a centralized authentication service can be reflected by recording an association between their respective service definitions. The recorded associations establish a service dependency topology that informs the data or selection options presented in a GUI, for example. (The service dependency topology is like a “map” showing how services are connected based on their dependencies.) The service topology may itself be depicted in a GUI and may be interactive to allow navigation among related services.

Entity definitions in SPLUNK® IT SERVICE INTELLIGENCE™ can include informational fields that can serve as metadata, implied data fields, or attributed data fields for the events identified by other aspects of the entity definition. Entity definitions in SPLUNK® IT SERVICE INTELLIGENCE™ can also be created and updated by an import of tabular data (as represented in a CSV, another delimited file, or a search query result set). The import may be GUI-mediated or processed using import parameters from a GUI-based import definition process. Entity definitions in SPLUNK® IT SERVICE INTELLIGENCE™ can also be associated with a service by means of a service definition rule. Processing the rule results in the matching entity definitions being associated with the service definition. The rule can be processed at creation time, and thereafter on a scheduled or on-demand basis. This allows dynamic, rule-based updates to the service definition.

During operation, SPLUNK® IT SERVICE INTELLIGENCE™ can recognize so-called “notable events” that may indicate a service performance problem or other situation of interest. These notable events can be recognized by a “correlation search” specifying trigger criteria for a notable event: every time KPI values satisfy the criteria, the application indicates a notable event. A severity level for the notable event may also be specified. Furthermore, when trigger criteria are satisfied, the correlation search may additionally or alternatively cause a service ticket to be created in an IT service management (ITSM) system, such as a systems available from ServiceNow, Inc., of Santa Clara, Calif.

SPLUNK® IT SERVICE INTELLIGENCE™ provides various visualizations built on its service-centric organization of event data and the KPI values generated and collected. Visualizations can be particularly useful for monitoring or investigating service performance. SPLUNK® IT SERVICE INTELLIGENCE™ provides a service monitoring interface suitable as the home page for ongoing IT service monitoring. The interface is appropriate for settings such as desktop use or for a wall-mounted display in a network operations center (NOC). The interface may prominently display a services health section with tiles for the aggregate KPI’s indicating overall health for defined services and a general KPI section with tiles for KPI’s related to individual service aspects. These tiles may display KPI information in a variety of ways, such as by being colored and ordered according to factors like the KPI state value. They also can be interactive and navigate to visualizations of more detailed KPI information.

SPLUNK® IT SERVICE INTELLIGENCE™ provides a service-monitoring dashboard visualization based on a user-defined template. The template can include user-selectable widgets of varying types and styles to display KPI information. The content and the appearance of widgets can respond dynamically to changing KPI information. The KPI widgets can appear in conjunction with a background image, user drawing objects, or other visual elements, that depict the IT operations environment, for example. The KPI widgets or

other GUI elements can be interactive so as to provide navigation to visualizations of more detailed KPI information.

SPLUNK® IT SERVICE INTELLIGENCE™ provides a visualization showing detailed time-series information for multiple KPI’s in parallel graph lanes. The length of each lane can correspond to a uniform time range, while the width of each lane may be automatically adjusted to fit the displayed KPI data. Data within each lane may be displayed in a user selectable style, such as a line, area, or bar chart. During operation a user may select a position in the time range of the graph lanes to activate lane inspection at that point in time. Lane inspection may display an indicator for the selected time across the graph lanes and display the KPI value associated with that point in time for each of the graph lanes. The visualization may also provide navigation to an interface for defining a correlation search, using information from the visualization to pre-populate the definition.

SPLUNK® IT SERVICE INTELLIGENCE™ provides a visualization for incident review showing detailed information for notable events. The incident review visualization may also show summary information for the notable events over a time frame, such as an indication of the number of notable events at each of a number of severity levels. The severity level display may be presented as a rainbow chart with the warmest color associated with the highest severity classification. The incident review visualization may also show summary information for the notable events over a time frame, such as the number of notable events occurring within segments of the time frame. The incident review visualization may display a list of notable events within the time frame ordered by any number of factors, such as time or severity. The selection of a particular notable event from the list may display detailed information about that notable event, including an identification of the correlation search that generated the notable event.

SPLUNK® IT SERVICE INTELLIGENCE™ provides pre-specified schemas for extracting relevant values from the different types of service-related event data. It also enables a user to define such schemas.

### 3.0 CODE COMPLETION

In some embodiments, the present disclosure provides for approaches to code completion that can be implemented to assist a user as the user codes commands of a coding language. Aspects of the present disclosure relate to various approaches to assisting a user in understanding the various coding options available to the user in a coding language, properly codifying those coding options in the coding language, as well understanding the functionality of the code they create. These approaches can speed up the time it takes a user to code, increase the likelihood that the code users create instructs a system to perform as intended, as well as reduces the number of times drafts of the code need to be executed and/or compiled by the system.

The coding language can be a query language, a pipelined coding language, such as SPL (e.g., a pipelined search language), or any other suitable coding language. A coding language typically allows a user to enter commands, which instruct the computer to perform tasks defined by the commands. In SPL, for example, these commands can include search commands that may operate on events, as described above.

Some coding languages allow a user to specify or provide a set of arguments (i.e., one or more arguments) and/or values of one or more arguments of a command. By pro-

viding different combinations of values and/or arguments for different instances of the same command, users are able to specify in code, different corresponding variations to how the tasks defined by the commands are to be performed.

To ensure a command is properly interpreted by a computer, the command should be written in proper syntax and include valid values for arguments, when present in the code. Syntax provides a set of rules that defines the combinations of symbols that are considered to be a correctly structured document or fragment in a coding language. For various reasons, a user may not provide one or more commands in proper syntax, may not be aware of the various commands available for coding in the system, or not be aware of various arguments and/or values thereof available for commands.

One such reason for the forgoing is the user could be a novice to coding or to a particular coding language. Furthermore, even experienced users can make errors in code, forget or be unaware of various coding options (e.g., obscure commands or arguments), and be unfamiliar with updates or changes to the coding language. This is compounded in cases where the coding language is complex and includes a large number of variations users could specify for commands.

Additionally, SPL allows users to define commands and syntax for the commands using a syntax coding language, meaning the available commands and variations thereof available to a particular user, who may not have coded the command, can be vast and can change over time.

In some cases, the search system described herein provides different applications that use the coding language. Different sets of commands with corresponding syntax can be provided to and used in different applications. A command can include a command identifier, which identifies the particular command from a set of commands in the coding language. In the present example, unbeknownst to a user, in different applications, different commands may use the same command identifier, but have different functionality and/or syntax.

### 3.1 Example Coding Interface

Aspects of various approaches to assisting users in coding are described below with respect to a screen of a coding interface illustrated by FIGS. 18A, 18B, 18C, 18D, 18E, 18F, 18G, 18H, 18I, 18J, 18K, 18L, 18M, and 18N (collectively referred to herein as "FIG. 18"). FIG. 18 depict a progression of the screen as the user enters code using the coding interface. The figures are depicted in sequence, such that FIG. 18A represents the screen at an initial point of the sequence, and FIG. 18N represents the screen at a final point of the sequence. However, any number of representations of the screen in-between the given figures may not be shown for conciseness of disclosure.

The present example depicted in FIG. 18 users SPL, as an example only. It should be appreciated that the present disclosure is not limited to SPL, or to the coding interface corresponding to FIG. 18. Additionally, it should be appreciated that at least one aspect of the disclosure implemented in a coding interface may not necessarily be implemented in another coding interface. In particular certain features may be independent from one another even though the coding interface shown may include a combination of those features.

In some implementations, the screen (e.g., a coding screen) of FIG. 18 corresponds to a search screen of a search system, such as search screen 600, described above. How-

ever, various aspect of the present disclosure need not be incorporated into a search screen of a search system, as shown. In some cases, the system provides an integrated development environment (IDE), which may not necessarily be incorporated into a search system.

In various implementations, the coding interface includes a text input region, which receives text input from the user. The text input region can display the text entered by the user, which can visual represent code of one or more commands being coded by the user using text. The text input regions can support word processing capabilities, such as any suitable combination of text input, deletion, modification, cursor position, visual indication of cursor position, copy, paste, the ability to set cursor position, mouse input, keyboard commands, and the like. Examples of potential features will be described below. In the present example, the text input region is implemented in search bar 1802 as a text box, which could be an HTML form. In other cases, the text input region could take other forms.

### 3.2 Contextual Suggestions And Syntax Templates

In various implementations, the system causes a set of one or more contextual suggestions to be displayed to the user based on input provided by the user to a text input region. Examples of such suggestions include suggestions 1806A, 1806B, 1806C, 1806D, 1806E, 1806F, 1806G, 1806H, and 1806I (also referred to collectively as suggestions 1806).

As an overview, the user can select a suggestion to cause code corresponding to the suggestion to be entered into the text input region. In some cases, the code is automatically entered as at least one complete command (e.g., as text in the text input region) in proper syntax for the coding language. In other cases, one or more placeholder blocks (e.g., argument block 1810) or other representation of one or more portions of the code is automatically entered into the text input region (e.g., using text and/or graphical interface objects that may include text). The user can interact with any placeholder blocks that may be included in the text input region to complete the code in the text input region.

An example of a placeholder block includes an argument block which corresponds to one or more arguments of a command. The user fills in (e.g., types in or selects suggestions for), or selects values for, the one or more placeholder blocks and the system ensures that the completed code incorporates the filled in code with proper syntax. In some cases, one or more placeholder blocks can be active and/or inactive. Code corresponding to active placeholder blocks may be edited using different combinations of user input and may be represented differently in the text input region compared to inactive placeholder blocks. For example, the code for inactive placeholder blocks may be represented and interacted with similar to code that does not correspond to an placeholder block (e.g., as plaintext with character by character keyboard navigation). For example, one or more placeholder blocks may be disabled in response to the user selectively submitting the one or more of the placeholder blocks in the text input region, thereby disabling the one or more placeholder blocks. Subsequently, the user may optionally re-enable or re-active at least one placeholder block for editing in an active state, which can include modifying selected values of an placeholder block, deleting the placeholder block, or adding an addition placeholder block to a set of placeholder blocks.

Optionally, when editing active placeholder blocks, such as argument blocks, the user may refrain from filling in one or more placeholder blocks, and the system completes the

code in the text input region without the corresponding code of those argument blocks (e.g., arguments) once submitted, were the code may be optional as defined by the syntax. Additionally, in some cases, the user can add and/or remove placeholder blocks and/or argument blocks from the text input region and the code is completed in the text input region according to the placeholder blocks.

In further respects, one or more placeholder blocks may be added to the text input region as a set of placeholder blocks that are part of a syntax template. A command syntax template is a type of syntax template that corresponds to one or more commands, but corresponds to a single command in FIG. 18, by way of example only. In some implementations, each command a user could enter into the text input region has a corresponding command syntax template. Further, each command syntax template comprises at least one placeholder block, such as at least one argument blocks, where each argument block represents a different argument of the command. As indicated above, at least one placeholder block may be optional, meaning it corresponds to code that is not required for the code of the syntax template to be written in proper syntax. Also, as indicated above, while the placeholder blocks are editable by the user, the user may add and/or remove placeholder blocks (e.g., optional argument blocks corresponding to optional arguments for the command) from the command syntax template. Furthermore, each placeholder block of a command syntax template may be submitted collectively, which disables or deactivates the placeholder blocks, and optionally re-enabled collectively, by submitting or re-enabling the command syntax template. As indicated above, when a command syntax template is submitted and/or disabled, the system may disable a set of user input methods for the command syntax template and/or placeholder blocks, such as navigation controls between placeholder blocks and/or the ability to receive and select suggestions for the placeholder blocks. Also indicated above, a different set of user input methods may be activated for the command syntax template and/or placeholder blocks when those components are submitted and/or disabled (e.g., in an inactive state).

An argument syntax template is a type of syntax template that corresponds to one or more arguments of a command. An argument syntax template can be embedded or nested within a command syntax template, and can provide similar functionality as a command syntax template, as further described below. Furthermore, in some implementations, an argument syntax template can be embedded or nested within another argument syntax template.

In the examples shown, suggestions are displayed as corresponding options in a menu (e.g., dropdown menu), or list. However, the suggestions can be in any suitable graphical interface element or combination or grouping of graphical interface elements. The options displayed to the user are selectable by the user. For example, FIG. 18A shows eight selectable options or suggestions in suggestions 1806A, each on a corresponding line and each having a corresponding suggestion type listed in the corresponding line (e.g., “Recently Used” and “Command”). Suggestion types will be further described below. Rather than displaying long suggestions (e.g., using wrapping), a suggestion that exceeds a threshold length (e.g., number of characters) may be truncated in the display, and the full text of suggestion could be revealed in response to the user hovering over the suggestion (e.g., in a tooltip) or based on other user input.

As mentioned above, each suggestion is selectable by the user to cause corresponding code to be entered into the text input region. In the examples shown, the user can first select

a suggestion from a set of suggestions displayed to the user, and subsequently select to incorporate code corresponding to the selected suggestion into the text input region. For example, starting from search bar 1802, a user can use keyboard keys, such as “Up Arrow” and “Down Arrow” keys, to navigate through the list of suggestions. The system causes the currently selected suggestion to be visually indicated to the user. In the example, shown, the selected suggestion is highlighted. For example, in FIG. 18B, the user has navigated to a command having a command identifier “timechart.”

The user may then provide further input to select the option corresponding to the suggestion. For example, by clicking the suggestion with a mouse or selecting the “Enter” or “Space” key or another selection key(s). When the user selects the displayed option, the code can automatically be entered into the text input region (which could be in the form of a syntax template or plain text). Using approaches described herein, the user can quickly navigate to and select a desired suggestion, causing a representation of code to be automatically entered into the text input region and input control of the user to be automatically returned to the text input region for additional coding. In some cases, one or more options could cause at least one additional graphical interface element to be displayed to the user (e.g., additional options and/or forms which is not shown in the present example), which the user can interact with to cause the code to be entered into the text input region.

In the example of FIG. 18B, the user selects the highlighted suggestion to cause corresponding code to automatically be entered into the text input region. Input control is automatically returned to the text input region, which in this example is to argument block 1810 in a syntax template, but in other cases could be to a plain text portion of the text input region. A comparison between the text input region in FIGS. 18B and 18C reveals text representing the code that has been added to the text input region (e.g., subsequent to the “i” in “timechart”) in the form of a syntax template.

Although in the example shown, at least some text is entered into the text input region, at least some of the code may be represented, at least partially using other means. For example, FIG. 18C includes text of a syntax template representing the code, but could use any combination of suitable visual indicator(s) and/or graphical interface objects to represent the code in the text input region. Thus, it should be appreciated that placeholder blocks and/or command syntax templates are examples of representations of the code.

Suggestions can be contextually suggested using rules, which map the user input to suggestions and define the content of the suggestions. In some implementations, a rule can include at least one extraction rule, such as a regex, to extract at least a portion of user input provided to the text input region. A rule can further include one or more tokens to assign to extracted text, as well as instructions for the rule that are executed when the input is mapped to the rule, such as mappings to other rules and defining which values and/or value types are included in the suggestions. As least some of these instructions may be based on the tokens assigned to the user input.

In some cases, suggestions are generated by autocompleting one or more portions of the user input. One or more of the rules can be mapped to the autocompleted input, such as by applying a regex of a rule to the autocompleted input. Different suggestions may correspond to different variations of the autocompleted input generated by an autocomplete algorithm.

Suggestions **1806A** and **1806B** are examples of suggestions that may be generated using autocomplete. In these examples, the set of suggestions is displayed based on at least one or more characters provided by the user to the text input region. Prior to FIG. **18A**, as an example, a user has typed all of the text into the text input region, up to the cursor (i.e., “sourcetype=access\_combined|”). When the user types in t, suggestions **1806A** automatically appear, as shown. The system has mapped the “|” character and the “t” character in the text input region to one or more rules, which define suggestions **1806A**. In the present example, “|” is referred to as a pipe symbol, which indicates a command can follow, according to the syntax of the coding language. Thus, it will be appreciated that the rules can be constructed to enforce at least some of the syntax of the coding language employed onto one or more of the suggestions.

Assume the user types “i” in the screen shown in FIG. **18A**. In the present example, this results in FIG. **18B** with suggestions **1806B** that are updated based on the additional input. Suggestions **1806B** can be generated similar to suggestions **1806A**, the output of the autocomplete algorithm being updated based on the additional input. As shown, the suggestions defined by the one or more rules include suggested command syntax templates and recently used commands.

A suggestion for a recently used command can correspond to a complete command in proper syntax. When a user selects a recently used command, the code can automatically be added to the text input region, such that the code in the text input region is executable by the system (e.g., the code may be incorporated into code (e.g., a query or portion thereof) already in the text input region). The recently used commands can be generated based on commands that the system has previously received in association with the user. For example, a recently used command can be suggested based on the system having performed a search in association with a user that included the command. The user may be the user currently constructing code in the text input region, or a user associated with that user, such as a user of a common client account, and/or having corresponding permissions or roles. A recently used command may be suggested based on having been previously run in association with the search interface, such as the interface of FIG. **18**, and/or another application where the execution is detectable by the system (e.g., via a search log).

In some cases, a recently used command may be suggested based on being previously entered into the text input region, such as in a previous coding session, or the same coding session after having been deleted from the text input region. Further, the code may have been entered into the text input region using any suitable input mechanism, such as via any suitable suggestion, as described herein. Another possible way a recently used command could be entered into text input region is by the user directly typing or pasting the command into the text input region.

As mentioned above, another example of a suggestion is a command syntax template. For example, suggestions **1806A** comprise five command syntax templates for commands having the command identifiers “table,” “top,” “tags,” “timechart,” and top. When a user selects to enter a command syntax template into the text input region, a set of one or more argument blocks of the command syntax template is added to the text input region.

Suggestions **1806A** and **1806B** are automatically displayed as a user types text into the text input region and correspond to the text typed by the user. However, contextual suggestions can be triggered in other ways and/or can

correspond to other portions of the text input region. Suggestions **1806C**, **1806D**, **1806E**, **1806F**, **1806G**, **1806H**, and **1806I** are examples of suggestions that correspond to placeholder blocks in the text input region. As mentioned above, each placeholder block represents one or more portions of code in the text input region. In the present examples, placeholder blocks will be described in the context of argument blocks, but the description can equally apply to placeholder blocks. In particular, each argument block represents one or more arguments of a command, but the description herein also applies to code that is not specifically an argument of a command. Furthermore, each argument block is part of a command syntax template that defines a single command. However, argument blocks may be used independent of command syntax templates and a command syntax template could correspond to multiple commands.

In the present examples, each argument block corresponds to a single argument of a command, by way of example only. Also, a single argument is visually indicated in the text input region by being bounded by “<” and “>.” The text input region can include any number of argument blocks, corresponding to different arguments of a command. For example, FIG. **18C** shows nine argument blocks, or placeholder blocks, in the text input region. In some cases, a user can navigate to an argument block in the text input region to select the argument block. A currently selected argument block(s) is visually indicated to the user. The example shown highlights selected argument blocks to identify the selected argument blocks to the user. The user can also deselect argument blocks, which may occur automatically based on selection of a previously unselected argument block. Optionally, when a command syntax template is added to the text input region, as described above, at least one argument block is automatically selected (e.g., the first argument block of the command syntax template). The argument block can be in an argument syntax template embedded in the command syntax template, which may include an argument block and optionally static code or text. In FIG. **18**, an argument syntax template is visually indicated as being bounded by “[” and “].”

In FIG. **18C**, for example, argument block **1810** has been automatically selected in response to the command syntax template being added to the text input region. In FIG. **18D**, the user has navigated from argument block **1810** (as shown in FIG. **18C**) to argument block **1812** (as shown in FIG. **18D**). In this example, each intermediate argument block has been navigated to in order to reach argument block **1812** (e.g., using a “Tab” key to navigate to the next argument block). Furthermore, navigating to an argument block causes the argument block to become selected and the previously selected argument block to be deselected. The user could similarly navigate in the other direction (e.g., using a “Shift-Tab” key combination), as desired, and navigation may wrap around for a currently activated command syntax template.

As shown suggestions **1806C**, **1806D**, **1806E**, **1806F**, **1806G**, **1806H**, and **1806I** are contextual based on the currently selected argument block(s). For example, suggestions **1806C** corresponds to argument block **1812**. As discussed previously, the aforementioned rules can be used to generate the suggestions for the selected argument block(s). In the example shown, the rules are based on the syntax of the command for the selected argument. For example, a rule can define one or more valid values of an argument block. In the example of FIG. **18D**, the rules are based on syntax for the “timechart” command and define valid values for the selected argument block based on various options that may be used as values for the argument in the command. In this

example, each option is defined by the syntax to include a static portion, that is to be used for each value corresponding to the option, as well as formatting for the option, and a variable portion. The syntax can also define one or more valid data types for the variable portion. As shown, a “span” option from the syntax corresponds to the selected suggestion “span=<log-span>” where “span=” is a static string followed by a variable portion of type log-span. The “span” option also corresponds to the non-selected suggestion “span=<span-length>,” based on the syntax defining a potential alternative typing for the variable portion as type span-length. The other options shown can be similarly defined and allow for any number of data types.

Thus, the system can generate the suggestions for an argument block based on the various options and variations thereof for the corresponding argument defined by the syntax. In the present example, the system is configured to provide a suggestion for each defined syntax variation of the argument in the command in suggestions **1806C**. In other cases, a subset of the syntax variations can be suggested. As another example, the same syntax variation can correspond to multiple suggestions, but the system may populate the variable portion with different values. In cases where a variable portion is populated, the system may determine the value(s) based on an analysis of the sourcetype (e.g., data, such as events in the sourcetype), previous search results associated with the user (e.g., returned by a search entered into the text input region), predefined rules of the corresponding data type (e.g., that specify a method and/or function to determine or calculate the value(s)), and/or other data.

In the present example, the user selects to enter code corresponding to the selected argument block into the text input region (e.g., similar to what has been described above with respect to FIGS. **18B** and **18C**). This causes text corresponding to the code to be automatically entered into the command syntax template at a location defined by the argument block (e.g., at the display location of the argument block). As shown in FIG. **18E**, the variable portion has been automatically filled in by the system and the code reads “span=1”. However, the variable portion could be blank, have placeholder text, or have a placeholder block, which as described above, may operate similar to an argument block (e.g., including offering suggestions). Entering the selection into the command syntax template optionally automatically returns input control to the argument block in a plain text editing mode, as shown. The user can input an “Enter” or “Tab” key to navigate to the next argument block in the command syntax template, if one is present.

FIGS. **18D** and **18E** illustrate implementations where the system can insert an additional placeholder block and/or argument syntax template in addition to the code corresponding to the selected argument block. In particular, the system has added argument block **1820** to the text input region. This behavior can be defined by the command syntax template and/or the argument syntax template corresponding to argument block **1814**. In the present example, the argument block is automatically added based on the syntax defining that at least one additional argument can be included in the argument syntax template. For example, the “. . .” in following an argument block in the text input region in FIG. **18** indicates to the user that at least one additional argument block can be added to the argument syntax template.

Using this approach, the user can add additional placeholder blocks and/or argument syntax templates to the argument syntax template and/or the command syntax tem-

plate when allowed for by the syntax for those templates. Thus, the user is informed as to whether additional options are available, without needing prior knowledge of the syntax, and can add additional code when desired without having to study extensive suggestions that list all of the possible combinations and variations of the code corresponding to the suggestions (which may be in the hundreds or more).

In the present example, the user can navigate to argument block **1820**, and may fill in a value for argument block **1820** to use the argument block in the completed code, or lead the argument block blank to complete the code without a corresponding argument. If the user fills in the value, it may be accomplished similar to argument block **1812**. Furthermore, an additional argument block could similarly be added if additional options are available (e.g., until the user exhausts the list of options defined by the syntax).

Similar to FIG. **18D**, suggestions **1806D** are based on argument block **1814**. The syntax of the command defines that a value of the argument can comprise one or more of the functions, as shown, which are captured by the rules applied to the argument block. As shown, at least one suggestion may correspond to an aggregator. An aggregator is a type of function that is always applied to each field of events in search results, as opposed to one or more fields when specified in the code.

It is also noted that in some implementations, the system refrains from displaying or providing one or more suggestions to the user based on the text input region already including code corresponding to the suggestion. In the present example, the system may refrain from including a suggestion corresponding to the suggestion selected in FIG. **18D** when displaying suggestions for argument block **1820** based on the corresponding suggestions being in the text input region (e.g., within the same argument syntax template or syntax template level).

Although suggestions may be presented to the user for an argument block, the user need not select one of the suggestions to enter a value(s) for the argument block. In FIG. **18F**, for example, assume the user types in “c,” rather than navigating into suggestions **1806D** (e.g., the dropdown menu). The user could then see updated suggestions based on autocompletion of the user input, which may only include the functions in suggestions **1806D** that begin with “c.” The user may then select the suggestion corresponding to “count( )” resulting in FIG. **18G**. As an alternative, the user could have typed in the corresponding code to result in FIG. **18G**.

It is noted that as a user types text into an argument block, suggestions are displayed based on the text input by the user for the value(s) of the argument block. These suggestions can be similar to what has been described above with respect to FIGS. **18A**, and **18B**. However, the suggestions are contextual to the argument block and the syntax defined for the argument block (e.g., for the argument syntax template). For example, functions may be suggested based on the entered text corresponding to a function identifier (e.g., count).

As an example, assume the user types “count(” in FIG. **18F**. When the user types in “(” the system can identify the user is currently entering in a value of a field for the function (e.g., using a regex that extracts the character from user input). This could also be based on the cursor position in the text input region. Based on the identification, the system can display suggestions corresponding to the variable portion of the syntax for the function and/or command. In the present example, where the user has selected the option corresponding to “count( )” suggestions could be generated a similar

way, or could be specific to the suggestion. In this case, the system determined fields for suggestions based on an analysis of previous search results corresponding to the sourcetype. The user can complete entering the value for argument block **1814** by inputting the “Enter” key or using some other manner of input, such as by selecting one of suggestions **1806E**, resulting in FIG. **18H**.

FIG. **18H** illustrates that a suggestion for an argument block (e.g., argument block **1822**) can optionally include one or more syntax templates. In this example, the user has selected an argument syntax template to fill in an argument block. However, other valid values include fields, as shown, or could include a placeholder block, depending on the syntax of the command and/or argument corresponding to the selected argument block.

Assume the user selects the selected suggestion of FIG. **18H**, to cause the argument syntax template to be incorporated into the text input region. In this case, the syntax template (e.g., the argument syntax template) becomes embedded or nested within the syntax template it was suggested for and may be referred to as an embedded syntax template or more specifically an embedded argument template in this example. The nested argument template includes at least one nested argument or placeholder block (e.g., argument block **1824**).

Any number of syntax templates can be embedded at a given layer of a syntax template and/or any number of layers of syntax templates may be embedded within the syntax template and/or placeholder block thereof depending on the definition of the syntax for the syntax template. Thus, a further syntax template could be nested in the syntax template corresponding to argument block

FIG. **18I** shows the resultant screen after the selection, and after the user has further filled in some of the values for nested argument blocks in the nested argument template using any suitable combination of approaches described above. Assume the user navigates to argument block **1824** causing suggestions **1824** to be displayed to the user, and selects the suggested argument syntax template, as shown. Further assume, the user selects to incorporate the argument syntax template corresponding to the selected suggestion to the text input region and fills in “count” into the first argument block that was added to the text input region. Suggestions **1806H** are displayed to the user, and the user navigates to a suggestion, resulting in FIG. **18J**.

FIGS. **18J** and **18K** illustrates a syntax history suggestion that can be included in the suggestions displayed to the user (or could be displayed elsewhere in the screen for selection and/or available for selection as a keyboard or other user input). The syntax history suggestion may be available to the user based on the suggestion being for a nested placeholder block and/or syntax template. In the present example, the suggestion is based on being for a nested argument block. Selection of a syntax history suggestion causes the system to automatically remove the nested code from the text input region. Furthermore, the nested code could be replaced with its parent code, such as its parent argument block, as shown in FIG. **18K**. Thus, the code in the text input region has returned to the historical state of FIG. **18I**, where the user could change their selected suggestion or enter text directly into argument block **1824**.

In various implementations, each syntax history suggestion is selectable to cause the code of the overall syntax template to return to a state a single level up from the currently selected placeholder block, by way of example. Furthermore, in various implementations, values entered into suggestions are retained for lower level placeholder

blocks. Thus, in FIG. **18K**, assume the user selects the same suggestion that was selected in FIG. **18I** to reintroduce the corresponding code into the text input region. As shown in FIG. **18L**, the value corresponding to “count” has been retained, and is automatically included in the code. The value has been stored in association with the suggestion despite the code having been removed from the text input region.

In the present example, the user subsequently enters a value include a remaining argument block, as shown in FIG. **18M**. The user then selectively submits the command syntax template (i.e., the base or outermost syntax template in this example), resulting in FIG. **18N**. In the implementation shown, the user can submit a base syntax template to the text input region in various ways. At any point, the user could press the “Escape” key. As another example, if the user is on the last argument block, such as in FIG. **18M**, pressing the “Enter” key submits the base syntax template. As mentioned previously, submitting a base syntax template can enable a set of user input methods for the code of the base syntax template and may disable a different set of user input methods utilized to edit the code while the base syntax template is active. The enabled set of user input methods can apply to other code in the text input region, such as code corresponding to other syntax templates and/or plain text in the text input region.

In some cases, when a base syntax template is disabled in the text input region, the user can edit the code of the base syntax template as plain text along with other plain text in the text input region. For example, in FIG. **18N**, all of the code of the base syntax template is represented as plain text, which the user can delete, add to, or otherwise manipulate using word processing input. Additionally, based on a user selecting to disable a base syntax template, the representation of the code in the text input region is modified in various ways. This includes, for example, removing syntax indicators which are not part of the code the user constructed, but indicate parts of the syntax to the user. Examples include argument or placeholder block indicators “<” “>,” embedded syntax template indicators “[” and “],” additional option indicators “. . .,” as well as embedded syntax templates that include blank, empty, or null values. Submitting a base syntax template can convert the representation into plain text that is in proper syntax for execution by the system, as shown.

The user could continue to code in the text input region shown in FIG. **18N**, which can include adding more base syntax templates to the text input region and/or typing in additional code. In some implementations, only a single base syntax template may be active at a given time. Furthermore, the user could be restricted from modifying code outside of the base syntax template in the text input region while the base syntax template is active. Also, in some cases, the user can selectively re-activate a base syntax template. For example, the user could type in a keyboard command or otherwise provide input to re-activate a base syntax template. Doing so in FIG. **18N**, for example, would reintroduce the representation of the syntax template that was previously in the text input region and the user can continue editing the code corresponding to the syntax template. The representation could look as it does in FIG. **18M**.

It is noted that in the present example, a user selects a syntax template as a suggestion from a list of suggestions. However, in other cases, the user could select any of the various syntax templates in a different manner. For example, a user could select a command syntax template from a list of command syntax templates, from a non-contextual menu, or

using any other suitable approach. In various implementations, by selecting and filling in syntax templates, the user can iteratively build commands, without needing to navigate comprehensive lists of options for a command by refraining from filling in optional placeholder blocks of a command syntax template, adding any number of placeholder blocks to the command syntax template, and/or removing any number of placeholder blocks from the command syntax template, as permitted by the syntax of the command.

### 3.3 Rule Generation

As described above, the system can use rules to dictate the suggestions provided to users when coding commands using the coding language. Other uses for these rules include data type validation. For example, using a regex of a rule, the system can extract a value provided by the user to the text input region (e.g., in a base syntax template or otherwise), and determine whether the value is of a valid data type. Whether a value is of a valid data type can be defined by the syntax of the command and/or coding language and incorporated into the rules. Based on identifying an invalid value, the system can visually indicate the value is invalid to the user. For example, the system can use the value to be highlighted or otherwise visually identified in the text input region, pop up an error message, or otherwise notify the user.

Another use for the rules includes syntax highlighting, which is shown in FIG. 18, and FIGS. 20A and 20B, by way of example. The syntax highlighting assists the user in visually identifying different elements of the syntax of the coding language on the screen. For example, a color coding system can be used to visually identify different elements of the syntax to the user. The rules are used to identify what characters in the text input region to highlight and what highlighting (e.g., color or formatting) to use in the text input region. In the example shown, highlighting is indicated using particular text formatting for common element types. To use the rules, the text is parsed to identify the corresponding parts or elements of the syntax. These elements can be represented as the tokens, described above, and the appropriate highlighting can be applied to the identified elements.

In some implementations, the system determines a respective set of rules for each command available to the user. For example, as described above, in some cases, different sets of commands may be available to different application and in some cases, commands can be defined individual using custom syntax. Thus, different parsing rules may be needed for each command. As a specific example, in some commands, BY is a reserved word but in other commands it is not a reserved word. Thus, the syntax code of commands may be parsed per-command to create the parsing rules for the commands. These parsing rules can then be used to parse the text in the text input region based on the commands in the string. These per-command parsing rules may be generated dynamically (not hard coded in client side code) which allows the rules to be applied for custom commands. For example, rules for commands of an application may be generated at load time of an application. Rules could be updated or generated each time the application is loaded or in some cases may be saved in association with an application and reused when loading the application.

In some implementations, a client displaying the application receives a file comprising a syntax definition for any given command (e.g., at load time of the application). These syntax definitions for each command are used to

build per-command parsing rules. FIG. 19A shows an example of a syntax definition for a command called "foo" (e.g., has a command identifier "foo"). FIG. 19B shows an example of rules for the foo command the system generated by parsing the syntax definition assigned to the foo command.

As shown in FIG. 19A, the syntax definition specifies clauses, which correspond to text surrounded by "[" and "]"." In creating rules, the system can identify each clause and each identified clause may be defined as syntax template in the rules. Furthermore, a clause in the syntax definition can reference one or more other clauses, which the system can identify and each identified clause may be used as a nested or embedded syntax template in the rules. These rules can be utilized to implement the various suggestions and behavior of syntax templates described above with respect to FIG. 18. The system can further identify available options and variations for the commands from the syntax definition. Additional rules may be used to perform designated actions based on the system applying the command specific rules to code corresponding to the command.

### 3.4 Code Reformatting

Additional aspects of the present disclosure relate to allowing a user to reformat code entered into a text input region. FIGS. 20A and 20B show search bar 2002, which can, as example, correspond to search bar 1802 or other text input regions described herein. Assume the user has entered text into a text input region, as shown in FIG. 20A. The user may have typed in the code and/or generated portions of the code using one or more syntax templates, as described above. In some implementations, the user can selectively reformat the code in the text input region. FIG. 20B shows an example of reformatting of the code in FIG. 20A. The system may cause display of the reformatted code in response to user input, such as one or more keyboard commands. Optionally, the system may subsequently cause display of the original code in response to further user input, such as one or more keyboard commands, or otherwise apply a different formatting template to the code.

In the present example, the system reformats the code by adding one or more new lines, tabs, and/or spaces to the code. For example, the system can identify each command in the code and insert the new formatting elements based on the identified commands. In some cases, the code is reformatted such that an identified command is on a corresponding line(s) of the code (e.g., by inserting a new line(s), when needed). In further respects, spacing of identified commands can be modified. As an example, spacing (e.g., at least one space character or tab) can be added to the code preceding an identified command (e.g., on its respective line(s)). The system can base the spacing on the level of the command in the code. In particular, commands may be nested within other commands. The spacing can indicate the level of the command such that the system is configured to ensure at least some different commands on a common nesting level have common spacing preceding the command, as shown. In the present example, these nested commands include nested searches, or sub-searches, and the level of a command is defined by its search level. In the example shown, each sub-search is defined by code surrounded by "[" and "]"." Identifying nested commands, as well as where commands being and/or end can be accomplished using the aforementioned parsing rules or other rules.

Also in some implementations, at least some of the formatting may be preserved when transitioning the code to



the predefined code format or template of FIG. 20A. In the example shown, new lines and tabs have been preserved for the “eval” command, as well as the arguments “foo=bar” and “foo=baz.” Had those formatting characters not been included in the code, the reformatted code might include each of those elements on a common line.

### 3.5 Additional Implementations

FIG. 21 presents a flowchart illustrating a method in accordance with the disclosed embodiments. Each block illustrating methods in accordance with FIG. 21, and other methods described herein, comprises a computing process that may be performed using any combination of hardware, firmware, and/or software. For instance, various functions may be carried out by a processor executing instructions stored in memory. The methods may also be embodied as computer-usable instructions stored on computer storage media. The methods may be provided by a standalone application, a service or hosted service (standalone or in combination with another hosted service), or a plug-in to another product, to name a few.

At block 2102, display is caused of a set of argument blocks of a command based on syntax of the command. For example, the system of FIG. 18 can cause a command syntax template of a command to be displayed in search bar 1802, which comprises argument blocks of the command, as shown in FIG. 18C. The command syntax template can be added based on any suitable user selection, but in the present example is based on selection of a corresponding suggestion. The command syntax template enables the user to access/input each argument and variation thereof permitted by the syntax definition of the command. However, initially, the command syntax template may be configured to display an argument block for each required and/or optional argument block at a top level of the command, as defined by the syntax definition.

At block 2104, modification is caused to the displayed set of argument blocks based on the syntax of the command. For example, the system can modify the set of argument blocks by adding argument blocks to the set, removing argument blocks from the set, and/or entering or changing values of the argument blocks. Each of these tasks can be managed by the system to ensure the modified set of argument blocks are in conformance with the syntax of the command and defined executable code. As an example, FIG. 18I could be a result of the modification.

At block 2106, the command is caused to be coded based on the set of argument blocks. For example, the as shown in FIG. 18N, the system converts the command syntax template into executable code that is in the text input region based on the values of the argument blocks. The user may than input an execution command to cause the query in the text input region to be executed by the system

FIG. 22 presents a flowchart illustrating a method in accordance with the disclosed embodiments. At block 2202, display is caused of an argument block. For example, based on a user selection of a suggestion as described above, argument block 1824 is displayed with a set of displayed argument blocks of a command.

At block 2204, display is caused of a nested argument block within the argument block. For example, the system can cause argument blocks to be nested within argument block 1824 based on a selection of a suggestion in suggestions 1806G.

At block 2206, display is caused of a value received from a user of the nested argument block. For example, the user

can enter “count” into one of the nested argument blocks as shown in FIG. 18J, resulting in the value being displayed in the text input region.

At block 2208, removal is caused of the displayed nested argument block from the argument block. For example, the user can select the syntax history option in suggestions 1806H resulting in FIG. 18K, which shows the argument block’s argument syntax template has been removed from the text input region.

At block 2210, the nested argument block is caused to be added back to the argument block with the value. For example, the user may select the suggestion again in suggestions 1806I, and the argument syntax template is added back to the text input region as shown in FIG. 8L. The argument block includes the value based on the user having inputted the value the previous time the argument syntax template was in the text input region.

From the foregoing, it will be seen that this invention is one well adapted to attain all the ends and objects set forth above, together with other advantages which are obvious and inherent to the system and method. It will be understood that certain features and subcombinations are of utility and may be employed without reference to other features and subcombinations. This is contemplated by and is within the scope of the claims.

The invention claimed is:

1. A computer-implemented method for causing coding of a command of a coding language, via a user interface, the method comprising:

in response to receiving, via the user interface, a first user selection that indicates the command, causing display of a syntax template for the command in a text input region of the user interface, wherein the displayed syntax template comprises a set of argument blocks including a first argument block that corresponds to a first argument for the command;

displaying, within a second region of the user interface that is separate from the text input region, a second argument block that is excluded from the set of argument blocks and corresponds to an available argument for the command;

in response to receiving, via the user interface, a second user selection that indicates the second argument block, updating the display of the syntax template to include the second argument block in the set of argument blocks, and wherein the updated display of the syntax template enables providing a value for the second argument; and

in response to receiving via the user interface, a third user selection that indicates the value for the second argument, causing the coding of the command, wherein the coded command is displayed within the text input region and includes the second argument having the received value.

2. The computer-implemented method of claim 1, wherein causing coding of the command comprises incorporating the coded command into existing code of the text input region, wherein the existing code includes a plurality of additional coded commands of the coding language.

3. The computer-implemented method of claim 1, wherein causing coding of the command comprises incorporating the coded command into a query defined in the coding language.

4. The computer-implemented method of claim 1, wherein the coding language is a pipelined coding language and the command is a pipeline command of the pipelined coding language.

51

5. The computer-implemented method of claim 1, further comprising causing:

execution of the coded command in a search query; and display of search results of the search query, wherein the text input region displays the coded command.

6. The computer-implemented method of claim 1, wherein the coded command displayed in the text input region comprises a subset of the set of argument blocks based on a value of an argument for a given argument block of the set of argument blocks having an empty value.

7. The computer-implemented method of claim 1, wherein the value for the second argument represents a nested argument of an argument of a parent argument block in the set of argument blocks.

8. The computer-implemented method of claim 1, further comprising:

in response to a fourth user selection that indicates the first argument block of the set of argument blocks, displaying a list of values for the first argument.

9. The computer-implemented method of claim 1, further comprising:

in response to receiving the value for the second argument, automatically causing a third argument block to be added to the set of argument blocks displayed in the text input region.

10. The computer-implemented method of claim 1, wherein causing coding of the command comprises converting the set of argument blocks into a plain text representation of the command and displaying the plain text representation of the command in the text entry region.

11. The computer-implemented method of claim 1, further comprising:

in response to receiving a fourth user selection an option for submitting the set of argument blocks to the command, disabling a set of user input methods for the set of argument blocks.

12. The computer-implemented method of claim 1, further comprising:

based on identifying the second argument block as a nested argument block of a parent argument block in a syntax of the command, causing display of a syntax history option that is selectable to cause removal of the second argument block from the set of argument blocks displayed in the text input region.

13. The computer-implemented method of claim 1, further comprising:

in response to a fourth user selection that indicates a syntax history option, causing removal of the second argument block from the set of argument blocks displayed in the text input region; and

in response to a fifth user selection that indicates the second argument block, causing the second argument block to be added back into the set of argument blocks with the value for the second argument.

14. The computer-implemented method of claim 1, further comprising:

parsing a syntax definition of the command; generating parsing rules from the parsed syntax definition of the command; and

causing the parsing rules to be applied to text, provided to the text input region, and generate a displayed list of values for the second argument.

15. The computer-implemented method of claim 1, further comprising:

52

causing receipt of a syntax definition for a plurality of commands of the coding language from a server in association with an application that provides the user interface;

causing command specific syntax rules to be generated from the syntax definition for each command of the plurality of commands; and

causing the application to enforce a set of the command specific syntax rules corresponding to the command on the set of argument blocks in the text entry region.

16. The computer-implemented method of claim 1, wherein the set of argument blocks are included in a command syntax template that enforces the syntax of the command on which argument blocks addable or removable from the set of argument blocks via an interaction with the command syntax template in the text input region.

17. The computer-implemented method of claim 1, wherein a third argument block of the set of argument blocks corresponds to an optional argument of the command, and the command is caused to be coded without the optional argument in the text input region based on being defined as optional in a syntax of the command.

18. The computer-implemented method of claim 1, causing a list of values for the second argument to be displayed, based on a given argument block of the set of argument blocks, and based on the given argument block being on a common level of a syntax with the second argument block.

19. One or more non-transitory computer-readable media having instructions stored thereon for causing coding of a command of a coding language, via a user interface, the instructions, when executed by a processor of a computing device, to cause the computing device to perform a method comprising:

in response to receiving, via the user interface, a first user selection that indicates the command, causing display of a syntax template for the command in a text input region of the user interface, wherein the displayed syntax template comprises a set of argument blocks including a first argument block that corresponds to a first argument for the command;

displaying, within a second region of the user interface that is separate from the text input region, a second argument block that is excluded from the set of argument blocks and corresponds to an available argument for the command;

in response to receiving, via the user interface, a second user selection that indicates the second argument block, updating the display of the syntax template to include the second argument block in the set of argument blocks, and wherein the updated display of the syntax template enables providing a value for the second argument; and

in response to receiving via the user interface, a third user selection that indicates the value for the second argument, causing the coding of the command, wherein the coded command is displayed within the text input region and includes the second argument having the received value.

20. The one or more computer-readable media of claim 19, wherein causing coding of the command comprises incorporating the coded command into a query defined in the coding language.

21. The one or more computer-readable media of claim 19, wherein the method further comprises causing:

execution of the coded command in a search query; and display of search results of the search query, wherein the text input region displays the coded command.

## 53

22. The one or more computer-readable media of claim 19, wherein the coded command displayed in the text input region comprises a subset of the set of argument blocks based on the value of an argument for a given argument block of the set of argument block having an empty value. 5

23. The one or more computer-readable media of claim 19, wherein the method further comprises:

in response to a fourth user selection that indicates the first argument block of the set of argument blocks, displaying a list of values for the first argument. 10

24. The one or more computer-readable media of claim 19, wherein the method further comprises:

in response to receiving the value for the second argument, automatically causing a third argument block to be added to the set of argument blocks displayed in the text input region. 15

25. A system comprising:

one or more processors; and

memory having instructions stored thereon for causing coding of a command of a coding language, via a user interface, the instructions, executable by the one or more processors to cause the system to perform a method comprising: 20

in response to receiving, via the user interface, a first user selection that indicates the command, causing display of a syntax template for the command in a text input region of the user interface, wherein the displayed syntax template comprises a set of argument blocks including a first argument block that corresponds to a first argument for the command; 25

displaying, within a second region of the user interface that is separate from the text input region, a second argument block that is excluded from the set of argument blocks and corresponds to an available argument for the command; 30

in response to receiving, via the user interface, a second user selection that indicates the second argument block, updating the display of the syntax template to include 35

## 54

the second argument block in the set of argument blocks, and wherein the updated display of the syntax template enables providing a value for the second argument; and

in response to receiving via the user interface, a third user selection that indicates the value for the second argument, causing the coding of the command, wherein the coded command is displayed within the text input region and includes the second argument having the received value.

26. The system of claim 25, wherein causing coding of the command comprises incorporating the coded command into a query defined in the coding language.

27. The system of claim 25, wherein the method further comprises causing:

execution of the coded command in a search query; and display of search results of the search query, wherein the text input region displays the coded command.

28. The system of claim 25, wherein the coded command displayed in the text input region comprises a subset of the set of argument blocks based on a value of an argument for a given argument block of the set of argument blocks having an empty value.

29. The system of claim 25, wherein the method further comprises:

in response to a fourth user selection that indicates the first argument block of the set of argument blocks, displaying a list of values for the first argument.

30. The system of claim 25, wherein the method further comprises:

in response to receiving the value for the second argument, automatically causing a third argument block to be added to the set of argument blocks displayed in the text input region.

\* \* \* \* \*